arXiv:1603.05229v1 [math.ST] 16 Mar 2016

PAC-Bayesian bounds for the Gram matrix and least squares regression with a random design.

Olivier Catoni

March 17, 2016

Abstract: The topics dicussed in this paper take their origin in the estimation of the Gram matrix $\mathbb{E}\left(XX^\top\right)$ of a random vector $X \in \mathbb{R}^d$ from a sample made of n independent copies $X_1, \dots, X_n$ of X. They comprise the estimation of the covariance matrix and the study of least squares regression with a random design. We propose four types of results, based on non-asymptotic PAC-Bayesian generalization bounds: a new robust estimator of the Gram matrix and of the covariance matrix, new results on the empirical Gram matrix $\frac{1}{n}\sum_{i=1}^n X_i X_i^\top$, new robust least squares estimators and new results on the ordinary least squares estimator, including its exact rate of convergence under polynomial moment assumptions.

Keywords: Gram matrix, covariance matrix, least squares regression with a random design, robust estimation, ordinary least squares estimator, PAC-Bayesian generalization bounds.

MSC2010: 62J10, 62J05, 62H20, 62F35, 15A52

Contents

1. A robust Gram matrix estimate 5
1.1. Definition of a new estimator 5
1.2. Generalization bounds 7
1.3. Estimation of the eigenvalues 8
2. The empirical Gram matrix estimate 9
3. Estimation of the covariance matrix 18
4. Least squares regression 22
4.1. From Gram matrix to least squares estimates 22
4.2. A robust least squares estimator 24

CNRS – CREST, UMR 9194, 3 avenue Pierre Larousse, 92240 Malakoff, France. e-mail: olivier.catoni@ensae.fr

4.3. Generalization bounds for the empirical risk minimizer 26
4.4. Some lower bound for the empirical risk minimizer . . . 34
4.5. Exact convergence rate for the empirical risk minimizer 36
A. Proof of Proposition 1.1 on page 7 42
B. Proof of Proposition 2.2 on page 11 55
C. Proof of Lemma 2.5 on page 16 56
D. Proof of Proposition 3.3 on page 19 57
E. Proof of lemma 3.4 on page 20 58
F. Proof of Proposition 4.9 on page 34 59
G. Proof of Proposition 4.10 on page 36 63
H. Proof of Proposition 4.12 on page 38 64
I. Obtaining a quadratic form. 67
J. Computation of the robust Gram matrix estimator . . 74
J.1. Computation of the robust estimator in a fixed direction. 74
J.2. Computation of a robust estimate of the Gram matrix 76
K. Some simulation 77

Introduction

Let us consider n independent copies $(X_1, \dots, X_n)$ of a random vector X whose probability distribution $\mathbb{P}$ belongs to $\mathcal{M}_+^1\big(\mathbb{R}^d\big)$, the set of probability measures on $\mathbb{R}^d$ (equiped with the Borel σ-algebra).

The topics discussed in this paper take their origin in the estimation of the Gram matrix $G = \mathbb{E}\big(XX^\top\big)$ and comprise the estimation of the covariance matrix $G - \mathbb{E}(X)\mathbb{E}(X)^\top$ and least squares regression with a random design.

We propose four things. A new robust estimator of G, new results on the empirical estimator

$$\overline{G} = \frac{1}{n}\sum_{i=1}^n X_i X_i^\top,$$

new robust least squares estimators and new results on the ordinary least squares estimator

$$\widehat{\theta} \in \arg\min_{\theta\in\mathbb{R}^d} \sum_{i=1}^{n} \big(Y_i - \langle \theta, X_i\rangle\big)^2,$$

where $(X_1, Y_1), \dots, (X_n, Y_n)$ are n independent copies of a couple $(X, Y) \in \mathbb{R}^d \times \mathbb{R}$ of random variables. In particular we give the exact rate of convergence of $R(\widehat{\theta}) - \inf_{\theta\in\mathbb{R}^d} R(\theta)$, where $R(\theta) = \mathbb{E}\big[\big(Y - \langle \theta, X\rangle\big)^2\big]$.

Technically, our approach is based on the estimation of the quadratic form

$$N(\theta) \stackrel{\text{def}}{=} \mathbb{E}\big(\langle \theta, X\rangle^2\big) = \theta^\top G \theta,$$

that computes the energy in direction θ. It can also be seen as the square of the norm defined by the Gram matrix G. Recovering G from N can be done through the polarization identity

$$\xi^\top G \theta = \frac{1}{4}\big[N(\xi+\theta) - N(\xi-\theta)\big] = \frac{1}{2}\big[N(\xi+\theta) - N(\xi) - N(\theta)\big], \qquad (1)$$

that gives as a special case

$$G_{i,j} = e_i^\top G e_j = \frac{1}{4}\big[N(e_i+e_j) - N(e_i-e_j)\big] = \frac{1}{2}\big[N(e_i+e_j) - N(e_i) - N(e_j)\big],$$

where e_i are the vectors of the canonical basis of $\mathbb{R}^d$.

Our purpose is to define and study robust estimators, whose estimation error $|N(\theta) - \widehat{N}(\theta)|$ can be bounded with a probability close to one jointly for all values of $\theta \in \mathbb{R}^d$ under weak polynomial moment assumptions. More specifically, we make a θ-dependent assumption on the variance of $\langle \theta, X\rangle^2$, that takes the form

$$\sup\Big\{\mathbb{E}\big(\langle \theta, X\rangle^4\big), \theta \in \mathbb{R}^d, \mathbb{E}\big(\langle \theta, X\rangle^2\big) \le 1\Big\} \le \kappa,$$

implying that

$$\mathbf{Var}\big(\langle \theta, X\rangle^2\big) \le \big(\kappa - 1\big)\mathbb{E}\big(\langle \theta, X\rangle^2\big)^2.$$

This kurtosis coefficient measures the heaviness of the tail of the distribution of $\langle \theta, X\rangle^2$. To give a point of comparison, when X is a Gaussian random vector, whatever its Gram matrix maybe, the above assumption is satisfied for $\kappa = 3$.

Based on this assumption, we define an estimator $\widehat{N}(\theta)$ and prove a θ-dependent uniform bound on the estimation error $|N(\theta) - \widehat{N}(\theta)|$. More precisely, instead of bounding merely

$$\sup\Big\{\big|N(\theta) - \widehat{N}(\theta)\big| \,:\, \theta \in \mathbb{R}^d, \|\theta\| \le 1\Big\},$$

we bound with a probability close to one

$$\sup_{\theta\in\mathbb{R}^d}\left|\frac{N(\theta)}{\widehat{N}(\theta)}-1\right|,$$

with the convention that $0/0=1$ and $z/0=+\infty$ when $z>0$. Remark that this type of bound implies that it is possible to estimate exactly the null space $\mathbf{Ker}(G)$ with a probability close to one.

This new estimator, built on the same principles as the robust mean estimator of [10], is interesting in at least two ways. First, it can be actually used to estimate G, with increased performances in some heavy tail situations and with mathematical guaranties taking the form of non-asymptotic convergence bounds under weak hypotheses.

Second, it can be compared with the empirical estimate $\overline{G}$ of G and used as a mathematical tool to prove new generalization bounds for $\overline{G}$, under various hypotheses, including polynomial moment assumptions.

The estimation of the Gram matrix has many interesting possible applications. The most obvious one is to derive robust alternatives to the classical principal component analysis based on $\overline{G}$. In this paper however, we will rather focus on least squares regression with a random design. We propose and study new stable least squares regression estimators, and also provide new bounds for the ordinaray least squares estimator. We already studied robust least squares regression in [5], but this time, as the reader will see, we come up with simplified estimators and tighter results. We also come up with interesting new results about the ordinary least squares estimator. In particular we give its exact rate of convergence under polynomial moment assumptions, including the case when the noise, defined as $Y-\langle\theta_*,X\rangle$, where $\theta_*\in\arg\min_{\theta\in\mathbb{R}^d}\mathbb{E}\big[\big(Y-\langle\theta,X\rangle\big)^2\big]$, is not independent from X. When independence is not assumed, this exact rate is quite interestingly not equal to $\mathbb{E}\big[\big(Y-\langle\theta_*,X\rangle\big)^2\big]d/n$ and can depart from it by an arbitrarily large or small factor, as we show on examples.

Let us close this introduction with some precisions on our use of the big $\mathbf{O}$ notation. In this paper we will always prove precise and fully explicit non-asymptotic bounds. Nevertheless, as these bounds are sometimes difficult to read, we will use the big $\mathbf{O}$ notation to give a representation of their order of magnitude. When we write $A=\mathbf{O}(B)$, where A and B are two expressions depending on parameters of the problem, we mean that there is a numerical constant c such that $A\le cB$. When we write $A=\mathbf{O}_{n\to\infty}(B)$, where n is the sample size, we mean that there is a numerical constant c such that $\limsup_{n\to\infty}A/B\le c$. Remark that $A=\mathbf{O}\big(B\big)$ implies that $A=\mathbf{O}_{n\to\infty}\big(B\big)$, but that the reverse implication is false in general. The

notation $A = \mathbf{O}_{n\to\infty}\big(B\big)$ means that B bounds the order of magnitude of the first order term of A seen as a function of the sample size, whereas $A = \mathbf{O}\big(B\big)$ means that B bounds the order of magnitude of A in all circumstances.

1. A robust Gram matrix estimate

1.1. Definition of a new estimator. Following the same route as in [10] in a more elaborate setting, we define some M-estimator of $N(\theta)$, the energy in direction θ, and derive for it non-asymptotic deviation bounds that are uniform with respect to θ. To do this, we need to introduce the influence function $\psi : \mathbb{R} \to \mathbb{R}$ defined as

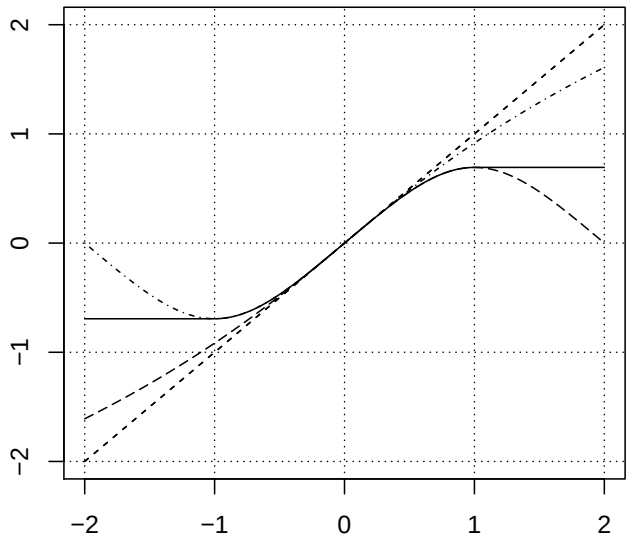

$x \mapsto \psi(x)$, compared with $x \mapsto x$
$x \mapsto \log\big(1 + x + x^2/2\big)$, and $x \mapsto -\log\big(1 - x + x^2/2\big)$

$$\psi(x) = \begin{cases} \log(2), & x \geq 1, \\ -\log\big(1 - x + x^2/2\big), & 0 \leq x \leq 1, \\ -\psi(-x), & x \leq 0. \end{cases}$$

This is a symmetric, non-decreasing, and bounded function, satisfying

$$-\log\big(1 - x + x^2/2\big) \leq \psi(x) \leq \log\big(1 + x + x^2/2\big), \qquad x \in \mathbb{R},$$

as can be seen from the identity

$$\big(1 - x + x^2/2\big)^{-1} = \frac{1 + x + x^2/2}{1 + x^4/4} \leq 1 + x + x^2/2, \qquad x \in \mathbb{R}.$$

Let us consider for any positive parameter $\lambda \in \mathbb{R}_+$ the empirical sample distribution

$$\overline{\mathbb{P}} = \frac{1}{n} \sum_{i=1}^{n} \delta_{X_i}$$

and the empirical criterion

$$r_\lambda(\theta) = \lambda^{-1} \int \psi\Big\{\lambda\big[\langle \theta, x\rangle^2 - 1\big]\Big\}\, \mathrm{d}\overline{\mathbb{P}}(x).$$

This empirical criterion depends on a positive scale parameter $\lambda \in]0, \infty[$, whose value will be set later on.

The centering of the influence function is done by adjusting the norm $\|\theta\|$, as we are going to explain now. To achieve centering, we introduce the multiplicative factor $\widehat{\alpha}(\theta)$ defined as

$$\widehat{\alpha}(\theta) = \sup\Big\{\alpha \in \mathbb{R}_+ \,:\, r_\lambda(\alpha\theta) \leq 0\Big\}.$$

As $r_\lambda(0) = -\lambda^{-1}\psi(\lambda) < 0$, we see that $\widehat{\alpha}(\theta) \in \mathbb{R}_+ \cup \{+\infty\}$, for any $\theta \in \mathbb{R}^d$. Moreover, since $\alpha \mapsto r_\lambda(\alpha\,\theta)$ is continuous, $r_\lambda\big(\widehat{\alpha}\,\theta\big) = 0$ as soon as $\widehat{\alpha}(\theta) < +\infty$. Considering that ψ is close to the identity on a neighbourhood of 0 and that the empirical measure $\overline{\mathbb{P}}$ is typically close to $\mathbb{P}$, we may hope that, for suitable values of λ, $r_\lambda(\theta) \simeq N(\theta) - 1$ with large probability. If this is the case, and if moreover, $\widehat{\alpha}(\theta) < \infty$, then

$$0 = r_\lambda\big[\widehat{\alpha}(\theta)\theta\big] \simeq N\big[\widehat{\alpha}(\theta)\theta\big] - 1 = \widehat{\alpha}(\theta)^2 N(\theta) - 1.$$

This is an incitation to define a new estimator of $N(\theta)$ as

$$\widehat{N}_\lambda(\theta) = \widehat{\alpha}(\theta)^{-2}. \tag{2}$$

To make things easier to understand, we can also define this estimator without introducing intermediate steps as

$$\widehat{N}_\lambda(\theta) = \inf\Big\{\rho \in \mathbb{R}_+^* \,:\, \sum_{i=1}^{n} \psi\Big[\lambda\Big(\rho^{-1}\langle \theta, X_i\rangle^2 - 1\Big)\Big] \leq 0\Big\}.$$

1.2. GENERALIZATION BOUNDS. We prove in Appendix A a detailed proposition, Proposition A.10 on page 53, whose main conclusions can be summarized as follows.

PROPOSITION 1.1 *Let us assume that for some known constant* κ

$$\sup\Big\{\mathbb{E}\big(\langle\theta,X\rangle^4\big)\,:\,\theta\in\mathbb{R}^d,\mathbb{E}\big(\langle\theta,X\rangle^2\big)\le 1\Big\}\le\kappa<\infty.$$

$$\textit{Choose } \lambda=\sqrt{\frac{2}{(\kappa-1)n}\big[\log(\epsilon^{-1})+0.73\,d\big]},$$
$$\textit{and consider } \mu=\sqrt{\frac{2(\kappa-1)}{n}\big[\log(\epsilon^{-1})+0.73\,d\big]}+6.81\sqrt{\frac{2\kappa\,d}{n}}.$$

For any confidence parameter $\epsilon>0$*, and any sample size* n *such that*

$$n>\left[20\sqrt{\kappa d}+\left(\frac{5}{2}+\frac{1}{2(\kappa-1)}\right)\sqrt{2(\kappa-1)\big[\log(\epsilon^{-1})+0.73d\big]}\right]^2 = \mathbf{O}\Big(\kappa\big[d+\log(\epsilon^{-1})\big]\Big), \quad (3)$$

with probability at least $1-2\epsilon$*, for any* $\theta\in\mathbb{R}^d$*,*

$$\left|\frac{N(\theta)}{\widehat{N}_\lambda(\theta)}-1\right|\le\frac{\mu}{1-2\mu}=\mathbf{O}\left(\sqrt{\frac{\kappa\big[d+\log(\epsilon^{-1})\big]}{n}}\,\right), \quad (4)$$

where it should be understood that $z/0=1$ *if* $z=0$ *and* $+\infty$ *if* $z>0$*, and where inequality* (4) *can be decomposed into two inequalities (by removing the absolute value) that hold each with probability at least* $1-\epsilon$*.*

The conclusion of this proposition is that the accuracy of the estimation of $N(\theta)$ by $\widehat{N}_\lambda(\theta)$ is of order $N(\theta)\sqrt{\kappa d/n}$ with a subgaussian tail up to very high (exponential with n) confidence levels. Indeed, equation (4) can also be written as

$$\Big|N(\theta)-\widehat{N}_\lambda(\theta)\Big|\le\frac{\mu\widehat{N}_\lambda(\theta)}{1-2\mu}\le\frac{\mu N(\theta)}{1-3\mu}.$$

This fairly strong result is made possible by the assumption that κ is bounded.

As already mentioned in the introduction, when X is a multidimensional Gaussian random variable, we can take $\kappa=3$. Assuming that $\kappa<\infty$ requires that the behaviour of the fourth moment of the distribution of $\langle\theta,X\rangle$ is not

too far from the Gaussian case. This remains nonetheless a much weaker assumption than the existence of exponential moments.

Let us mention that instead of getting a generalization bound or order

$$\mathbf{O}\left(\sqrt{\frac{\kappa\big[d+\log(\epsilon^{-1})\big]}{n}}\right),$$

depending on the dimension d of the ambient space (or more accurately on the rank of the Gram matrix G), it is also possible to obtain dimension free bounds where the dimension d is replaced with the θ dependent term $\mathbf{T}\mathrm{r}(G)\|\theta\|^2/N(\theta)$ (to see that it is indeed some substitute for the dimension, we can remark that this θ dependent factor is uniformly equal to the dimension d in the case when $G=\mathbf{I}$, because then $\mathbf{T}\mathrm{r}(G)=d$ and $N(\theta)=\|\theta\|^2$). For such results, we refer to the works of our student Ilaria Giulini [12, 13].

The estimator $\widehat{N}_\lambda(\theta)$ is not a quadratic form in θ, and therefore does not define an estimate of the Gram matrix G in an obvious way. Nevertheless, we show in appendix that it is possible to deduce from $\widehat{N}_\lambda$ a robust estimate $\widehat{G}$ of G.

Proposition 1.2 *There exists an estimator $\widehat{G}$ of the Gram matrix G, deduced from $\widehat{N}_\lambda$ as explained in Appendix I, such that under the same hypotheses as in Proposition 1.1, for any confidence parameter $\epsilon>0$ and any sample size n satisfying equation* (3)*, with probability at least $1-2\epsilon$, for any $\theta\in\mathbb{R}^d$,*

$$\left|\frac{\theta^\top\widehat{G}\theta}{\theta^\top G\theta}-1\right|\leq\frac{2\mu}{1-4\mu}=\mathbf{O}\left(\sqrt{\frac{\kappa\big[d+\log(\epsilon^{-1})\big]}{n}}\right), \tag{5}$$

where μ is defined as in Proposition 1.1.

Proof. This proposition is a simplified formulation of the end of Corollary I.2 on page 70, where the construction is applied to $\widehat{N}_\lambda$ that satisfies equation (32) on page 67 with $\delta=\mu/(1-2\mu)$, so that

$$\frac{2\delta}{1-2\delta}=\frac{2\mu}{1-4\mu}.$$

□

1.3. Estimation of the eigenvalues. Let us mention that the result stated in Proposition 1.2 induces an estimation of the eigenvalues of G. Indeed, if $\lambda_1\geq\lambda_2\geq\cdots\geq\lambda_d$ are the eigenvalues of G (counted with their

multiplicities), and $\widehat{\lambda}_1 \geq \widehat{\lambda}_2 \geq \cdots \geq \widehat{\lambda}_d$ are the eigenvalues of $\widehat{G}$, then when equation (5) holds,

$$\sup_{i\in\{1,\dots,d\}} \left| \frac{\widehat{\lambda}_i}{\lambda_i} - 1 \right| \leq \frac{2\mu}{1-4\mu}.$$

Proof. Let $\mathbf{Gr}\big(\mathbb{R}^d, i\big)$ be the set of linear subspaces of $\mathbb{R}^d$ of dimension i. The above inequality is a direct consequence of the fact that

$$\widehat{\lambda}_i = \sup\left\{ \inf\Big\{ \theta^\top \widehat{G}\theta \;:\; \theta \in V \cap \mathbb{S}_d \Big\} \;:\; V \in \mathbf{Gr}\big(\mathbb{R}^d, i\big) \right\},$$

$$\text{whereas } \lambda_i = \sup\left\{ \inf\Big\{ \theta^\top G\theta \;:\; \theta \in V \cap \mathbb{S}_d \Big\} \;:\; V \in \mathbf{Gr}\big(\mathbb{R}^d, i\big) \right\}.$$

These two identities themselves can be established from the remark that for any $V \in \mathbf{Gr}\big(\mathbb{R}^d, i\big)$, and any orthonormal basis $(e_1, \dots, e_d)$,

$$\dim\Big(V \cap \mathbf{span}\big\{ e_i, e_{i+1}, \dots, e_d \big\} \Big) \geq 1,$$

considering the case when $(e_1, \dots, e_d)$ is a basis of eigenvectors of $\widehat{G}$ or of G, corresponding to the eigenvalues $(\widehat{\lambda}_1, \dots, \widehat{\lambda}_d)$ or $(\lambda_1, \dots, \lambda_d)$ respectively. □

2. The empirical Gram matrix estimate

In this section, we study the empirical Gram matrix estimate

$$\overline{G} = \frac{1}{n} \sum_{i=1}^{n} X_i X_i^\top,$$

and the corresponding quadratic form $\overline{N}(\theta) = \theta^\top \overline{G} \theta$. We use the previous robust estimate $\widehat{N}_\lambda(\theta)$ of $\theta^\top G \theta$ as a tool. As we will always use the value of λ defined in Proposition 1.1 on page 7, we will write in this section for short $\widehat{N}$ instead of $\widehat{N}_\lambda$. Our approach is to analyze the difference $\overline{N}(\theta) - \widehat{N}(\theta)$, showing that it is small under suitable assumptions.

First of all, we deduce from the definitions of $\widehat{N}$ and $\overline{N}$ that $\widehat{N}(\theta) = \overline{N}(\theta) = 0$ almost surely for any $\theta \in \mathbf{Ker}\, G$ (that is any θ such that $N(\theta) = 0$).

On the other hand, under the hypotheses of Proposition 1.1 on page 7, with probability at least $1-\epsilon$, $N(\theta)/\widehat{N}(\theta) \leq 1 + \widehat{\delta}$, so that in the case when $N(\theta) > 0$, $\widehat{N}(\theta) > 0$ also, so that $r_\lambda\big(\widehat{N}(\theta)^{-1/2}\theta\big) = 0$. As a consequence,

$$\frac{\overline{N}(\theta)}{\widehat{N}(\theta)} - 1 = \lambda^{-1} \int \lambda\big[\langle \theta, x \rangle^2 \widehat{N}(\theta)^{-1} - 1 \big] \, \mathrm{d}\overline{\mathbb{P}}(x)$$

$$= \lambda^{-1} \int g\big[\lambda\big(\langle\theta, x\rangle^2 \widehat{N}(\theta)^{-1} - 1\big)\big] \, \mathrm{d}\overline{\mathbb{P}}(x),$$

where $g(z) = z - \psi(z)$. It is easy to compute

$$g'(z) = \begin{cases} 1, & z \geq 1 \\ \dfrac{z^2}{1 + (z-1)^2}, & 0 \leq z \leq 1, \\ g'(-z), & z \leq 0, \end{cases}$$

showing that $g'(z) \leq z^2$, and therefore that $g(z) \leq \max\{z, 0\}^3/3$, for any $z \in \mathbb{R}$. We see also that for any $p \in [0, 2]$, and any $z \in \mathbb{R}_+$, $g'(z) \leq z^p$, so that more generally $g(z) \leq \max\{z, 0\}^{p+1}/(p+1)$ for any $z \in \mathbb{R}$ and any $p \in [0, 2]$. As a consequence

Proposition 2.1 *Let us make the same assumptions as in Proposition 1.1 on page 7. On an event of probability at least $1 - \epsilon$ that includes the event of probability at least $1 - 2\epsilon$ described in Proposition 1.1 on page 7,*

$$\begin{aligned} \frac{\overline{N}(\theta)}{\widehat{N}(\theta)} - 1 &\leq \inf_{p \in [0,2]} \frac{\lambda^p}{p+1} \int \Big(\langle\theta, x\rangle^2 \widehat{N}(\theta)^{-1} - 1\Big)_+^{p+1} \, \mathrm{d}\overline{\mathbb{P}}(x) \\ &\leq \frac{\lambda^2}{3} \int \left(\langle\theta, x\rangle^2 \widehat{N}(\theta)^{-1} - 1\right)_+^3 \mathrm{d}\overline{\mathbb{P}}(x), \end{aligned}$$

and

$$1 - \frac{\overline{N}(\theta)}{\widehat{N}(\theta)} \leq \frac{\lambda^2}{3} \int \Big(1 - \langle\theta, x\rangle^2 \widehat{N}(\theta)^{-1}\Big)_+^3 \mathrm{d}\overline{\mathbb{P}}(x) \leq \frac{\lambda^2}{3},$$

where $(z)_+ = \max\{z, 0\}$ and λ is defined as in Proposition 1.1 on page 7, so that

$$\lambda^2 = \frac{2\big[\log(\epsilon^{-1}) + 0.73\, d\big]}{(\kappa - 1)n}.$$

This proposition uses random upper bounds. Nevertheless, it gives an indication that in good cases, when the fluctuations of these random upper bounds are not too wild, the difference between $\overline{N}$ and $\widehat{N}$, measured by $\left|\dfrac{\overline{N}(\theta)}{\widehat{N}(\theta)} - 1\right|$, should be of order λ^2, that is of order n^{-1}, whereas, as we have already seen, $\left|\dfrac{N(\theta)}{\widehat{N}(\theta)} - 1\right|$ is of order $n^{-1/2}$. More precisely, the second inequality proves that $\overline{N}(\theta)$ cannot be significantly *smaller* than $\widehat{N}(\theta)$, while the first inequality shows that it can be significantly larger, but only in the case when the

fluctuations of the random quantity

$$\int \langle \theta, x\rangle^6 N(\theta)^{-3} \mathrm{d}\overline{\mathbb{P}}(x)$$

are not bounded with n, since with probability $1-2\epsilon$

$$\begin{aligned}\int \Big(\langle \theta, x\rangle^2 \widehat{N}(\theta)^{-1} - 1\Big)_+^3 \mathrm{d}\overline{\mathbb{P}}(x) &\leq \int \left(\langle \theta, x\rangle^2 N(\theta)^{-1}\left(\frac{1-\mu}{1-2\mu}\right) - 1\right)_+^3 \mathrm{d}\overline{\mathbb{P}}(x) \\ &\leq \left(\frac{1-\mu}{1-2\mu}\right)^3 \int \langle \theta, x\rangle^6 N(\theta)^{-3}\, \mathrm{d}\overline{\mathbb{P}}(x),\end{aligned}$$

where μ is defined as in Proposition 1.1 on page 7.

We will now replace the bounds in the previous proposition by more explicit ones.
Write the Gram matrix G in diagonal form as

$$G = U\, \mathbf{diag}\big(\lambda_1, \ldots, \lambda_d\big)\, U^\top,$$

where $UU^\top = \mathbf{I}$ and $\lambda_1 \geq \cdots \geq \lambda_d$, and define

$$G^{-1/2} = U\, \mathbf{diag}\Big[\mathbb{1}\big(\lambda_i > 0\big)\lambda_i^{-1/2}, i = 1, \ldots, d\Big] U^\top.$$

As almost surely $X_i \in \mathbf{Im}(G)$, almost surely $G^{1/2}\, G^{-1/2} X_i = X_i$ and therefore

$$\begin{aligned}\big\|G^{-1/2} X_i\big\| &= \sup\Big\{\big\langle G^{-1/2} X_i, \theta\big\rangle \,:\, \theta \in \mathbf{Im}(G), \|\theta\| \leq 1\Big\} \\ &= \sup\Big\{\big\langle G^{-1/2} X_i, G^{1/2}\theta\big\rangle \,:\, \theta \in \mathbb{R}^d, \big\|G^{1/2}\theta\big\| \leq 1\Big\} \\ &= \sup\Big\{\langle X_i, \theta\rangle \,:\, \theta \in \mathbb{R}^d, \mathbb{E}\big(\langle \theta, X_i\rangle^2\big) \leq 1\Big\}.\end{aligned}$$

Consider

$$R = \max_{i=1,\ldots,n} \|G^{-1/2} X_i\| = \max_{i=1,\ldots,n} \sup\Big\{\langle \theta, X_i\rangle \,:\, \theta \in \mathbb{R}^d, \mathbb{E}\big(\langle \theta, X\rangle^2\big) \leq 1\Big\}. \quad (6)$$

Using these remarks and this definition, we can state the following consequence of Proposition 2.1:

Proposition 2.2 *Define the quantities*

$$\widehat{\delta} = \frac{\mu}{1-2\mu}, \qquad \textit{where } \mu \textit{ is as in Proposition 1.1 on page 7,}$$

$$\gamma_+ = \frac{2\big[\log(\epsilon^{-1}) + 0.73\,d\big]\,R^4\,(1+\widehat{\delta})^2}{3(\kappa-1)n},$$

where R is defined in equation (6),

$$\gamma_- = \frac{2\big[\log(\epsilon^{-1}) + 0.73\,d\big]}{3(\kappa-1)n},$$

$$\begin{aligned}\overline{\delta}_+(\theta) &= \frac{2\big[\log(\epsilon^{-1}) + 0.73\,d\big]}{3(\kappa-1)n}\int\Big(\langle\theta,x\rangle^2\widehat{N}(\theta)^{-1}-1\Big)_+^3\,\mathrm{d}\overline{\mathbb{P}}(x)\\ &\le \gamma_+\frac{\overline{N}(\theta)}{\widehat{N}(\theta)}\\ \overline{\delta}_-(\theta) &= \frac{2\big[\log(\epsilon^{-1}) + 0.73\,d\big]}{3(\kappa-1)n}\int\Big(1-\langle\theta,x\rangle^2\widehat{N}(\theta)^{-1}\Big)_+^3\,\mathrm{d}\overline{\mathbb{P}}(x)\\ &\le \gamma_-.\end{aligned}$$

Under the hypotheses of Proposition 1.1 on page 7, using the above notation and definitions, with probability at least $1-2\epsilon$, for any $\theta\in\mathbb{R}^d$

$$\begin{aligned}-\underbrace{\frac{\widehat{\delta}+\gamma_-}{1+\widehat{\delta}}}_{=\mathbf{O}(\widehat{\delta})} \le -\frac{\widehat{\delta}+\overline{\delta}_-(\theta)}{1+\widehat{\delta}} \le \frac{\overline{N}(\theta)}{N(\theta)} - 1 \le \frac{\widehat{\delta}+\overline{\delta}_+(\theta)}{1-\widehat{\delta}}\\ \le \frac{1}{(1-\widehat{\delta})(1-\gamma_+)} - 1 \le \frac{\widehat{\delta}+\gamma_+}{(1-\widehat{\delta})(1-\gamma_+)},\end{aligned}$$

where it is useful to remember that $\widehat{\delta} = \mathbf{O}\left(\sqrt{\frac{\kappa[d+\log(\epsilon^{-1})]}{n}}\right)$.

The proof of this proposition is given in appendix. Let us remark that our lower bound for $\overline{N}(\theta)/N(\theta)$, that is always $\mathbf{O}\big(\widehat{\delta}\big)$ and holds under the hypotheses stated in Proposition 1.1 on page 7, can be compared with the lower bound on the smallest singular value of a random matrix with i.i.d. isotropic columns given in [14, Theorem 1.3]. The first point in this theorem of V. Koltchinskii and S. Mendelson gives a slightly worse lower bound with less explicit constants (in particular, the dependence in κ is not explicit) under a slightly stronger condition, whereas points 2. and 3. of their theorem prove slower rates than $1/\sqrt{n}$ under weaker assumptions than ours.

Let us now upper-bound the random quantity R defined by equation (6) and consequently γ_+ under suitable assumptions. A simple choice is to assume an exponential moment of the type

$$\mathbb{E}\left\{\exp\left[\frac{\alpha}{2}\Big(\|G^{-1/2}X\|^2 - d - \eta\Big)\right]\right\} \le 1, \tag{7}$$

where α and η are two positive real constants.

Under this assumption, with probability at least $1-\epsilon$,

$$R^2 \leq d + \eta + \frac{2}{\alpha}\log\bigl(n/\epsilon\bigr).$$

Remark that $\mathbb{E}\bigl(\|G^{-1/2}X\|^2\bigr) = d$, so that assumption (7) can also be written as

$$\mathbb{E}\biggl\{\exp\biggl[\frac{\alpha}{2}\Bigl(\|G^{-1/2}X\|^2 - \mathbb{E}\bigl(\|G^{-1/2}X\|^2\bigr) - \eta\Bigr)\biggr]\biggr\} \leq 1.$$

Let us also remark that in the case when $X \in \mathbb{R}^d$ is a centered Gaussian vector,

$$\mathbb{E}\biggl\{\exp\biggl[\frac{\alpha}{2}\Bigl(\|G^{-1/2}X\|^2 + \frac{d}{\alpha}\log(1-\alpha)\Bigr)\biggr]\biggr\} = 1, \qquad 0 < \alpha < 1,$$

so that we can take in this case $\eta = -\dfrac{d}{\alpha}\bigl[\log(1-\alpha) + \alpha\bigr]$. In the case when X is a non centered Gaussian vector, we can also check that

$$\mathbb{E}\biggl\{\exp\biggl[\frac{\alpha}{2}\Bigl(\|G^{-1/2}X\|^2 + \frac{d}{\alpha}\log(1-\alpha)\Bigr)\biggr]\biggr\} \leq 1, \qquad 0 < \alpha < 1,$$

so that the same choice of η is still valid.

Hypothesis (7) is quite strong, and can be replaced by the more general assumption that

$$\mathbb{E}\biggl\{\exp\biggl[\frac{\alpha}{2}\Bigl(\|G^{-1/2}X\|^{2p} - \mathbb{E}\bigl(\|G^{-1/2}X\|^{2p}\bigr) - \eta\Bigr)\biggr]\biggr\} \leq 1, \tag{8}$$

for some exponent $p \in]0,1]$ and positive constants α and η. Under this new assumption, with probability at least $1-\epsilon$,

$$\begin{aligned} R^2 &\leq \biggl(\mathbb{E}\bigl(\|G^{-1/2}X\|^{2p}\bigr) + \frac{2}{\alpha}\log\bigl(n/\epsilon\bigr) + \eta\biggr)^{1/p} \\ &\leq \biggl(d^p + \frac{2}{\alpha}\log\bigl(n/\epsilon\bigr) + \eta\biggr)^{1/p}, \end{aligned}$$

since

$$\mathbb{E}\bigl(\|G^{-1/2}X\|^{2p}\bigr) \leq \mathbb{E}\bigl(\|G^{-1/2}X\|^2\bigr)^p = d^p.$$

PROPOSITION 2.3 *Let us assume that condition* (8) *as well as the hypotheses of Proposition 1.1 on page 7 are satisfied, and introduce the constant*

$$\begin{aligned}\widehat{\gamma}_+ &= \frac{2\big[\log(\epsilon^{-1}) + 0.73\,d\big]\Big[\mathbb{E}\big(\|G^{-1/2}X\|^{2p}\big) + 2\alpha^{-1}\log(n/\epsilon) + \eta\Big]^{2/p}(1+\widehat{\delta})^2}{3(\kappa-1)n} \\ &\le \frac{2\big[\log(\epsilon^{-1}) + 0.73\,d\big]\big[d^p + 2\alpha^{-1}\log(n/\epsilon) + \eta\big]^{2/p}(1+\widehat{\delta})^2}{3(\kappa-1)n}.\end{aligned}$$

With probability at least $1-\epsilon$,

$$\gamma_+ \le \widehat{\gamma}_+,$$

so that with probability at least $1-3\epsilon$, *for any* $\theta \in \mathbb{R}^d$,

$$-\underbrace{\frac{\widehat{\delta}+\gamma_-}{1+\widehat{\delta}}}_{=\mathbf{O}(\widehat{\delta})} \le \frac{\overline{N}(\theta)}{N(\theta)} - 1 \le \underbrace{\frac{\widehat{\delta}+\widehat{\gamma}_+}{(1-\widehat{\delta})(1-\widehat{\gamma}_+)_+}}_{=\mathbf{O}_{n\to\infty}(\widehat{\delta})},$$

where the notations are the same as in Proposition 2.2 on page 11 and where it is useful to remember that

$$\widehat{\delta} = \mathbf{O}\left(\sqrt{\frac{\kappa\big[d+\log(\epsilon^{-1})\big]}{n}}\,\right).$$

We can also replace hypothesis (8) by a polynomial moment assumption. Remembering Proposition 2.1 on page 10, remark that on the set Ω of probability at least $1-2\epsilon$ appearing in Proposition 1.1 on page 7, for any $\theta \in \mathbb{R}^d$,

$$\begin{aligned}\frac{\overline{N}(\theta)}{\widehat{N}(\theta)} - 1 &\le \frac{\lambda^p}{p+1}\int\Big(\langle\theta,x\rangle^2\widehat{N}(\theta)^{-1} - 1\Big)_+^{p+1}\,\mathrm{d}\overline{\mathbb{P}}(x) \\ &\le \frac{\lambda^p}{p+1}\int\Big(N(\theta)\|G^{-1/2}x\|^2\widehat{N}(\theta)^{-1} - 1\Big)_+^{p+1}\,\mathrm{d}\overline{\mathbb{P}}(x). \quad (9)\end{aligned}$$

Since $N(\theta)\widehat{N}(\theta)^{-1} \le 1+\widehat{\delta}$ on the event Ω,

$$\int\Big(\langle\theta,x\rangle^2\widehat{N}(\theta)^{-1} - 1\Big)_+^{p+1}\,\mathrm{d}\overline{\mathbb{P}}(x) \le \int f(x)\overline{\mathbb{P}}(x),$$

where

$$f(x) = \Big((1+\widehat{\delta})\|G^{-1/2}x\|^2 - 1\Big)_+^{p+1}.$$

From Bienaymé Chebyshev's inequality, with probability at least $1-\epsilon$,

$$\int f(x)\,\mathrm{d}\overline{\mathbb{P}}(x) \leq \mathbb{E}\big[f(X)\big] + \left(\frac{\mathbf{Var}\big[f(X)\big]}{n\epsilon}\right)^{1/2} \leq \mathbb{E}\big[f(X)\big] + \left(\frac{\mathbb{E}\big[f(X)^2\big]}{n\epsilon}\right)^{1/2}.$$

This leads to

PROPOSITION 2.4 *Consider some exponent $p \in]1,2]$ and introduce the bound*

$$\begin{aligned}\widetilde{\gamma}_+ &= \frac{1}{p+1}\left(\frac{2\big[\log(\epsilon^{-1}) + 0.73\,d\big]}{(\kappa-1)n}\right)^{p/2}\Bigg[\mathbb{E}\Big[\Big((1+\widehat{\delta})\|G^{-1/2}X\|^2 - 1\Big)_+^{p+1}\Big] \\ &\qquad + \left(\frac{\mathbb{E}\big[\big((1+\widehat{\delta})\|G^{-1/2}X\|^2 - 1\big)_+^{2p+2}\big]}{n\epsilon}\right)^{1/2}\Bigg] \\ &\leq \frac{1}{p+1}\left(\frac{2\big[\log(\epsilon^{-1}) + 0.73\,d\big]}{(\kappa-1)n}\right)^{p/2}(1+\widehat{\delta})^{p+1}\Bigg[\mathbb{E}\big(\|G^{-1/2}X\|^{2p+2}\big) \\ &\qquad + \left(\frac{\mathbb{E}\big(\|G^{-1/2}X\|^{4p+4}\big)}{n\epsilon}\right)^{1/2}\Bigg].\end{aligned}$$

Under the hypotheses of Proposition 1.1 on page 7, with probability at least $1-3\epsilon$, for any $\theta \in \mathbb{R}^d$,

$$-\underbrace{\frac{\widehat{\delta}+\gamma_-}{1+\widehat{\delta}}}_{=\mathbf{O}(\widehat{\delta})} \leq \frac{\overline{N}(\theta)}{N(\theta)} - 1 \leq \underbrace{\frac{\widehat{\delta}+\widetilde{\gamma}_+}{(1-\widehat{\delta})(1-\widetilde{\gamma}_+)_+}}_{=\mathbf{O}_{n\to\infty}(\widehat{\delta})},$$

where the notations are otherwise the same as in Proposition 2.2 on page 11 and whese we recall that

$$\widehat{\delta} = \mathbf{O}\left(\sqrt{\frac{\kappa\big[d+\log(\epsilon^{-1})\big]}{n}}\right).$$

To give an idea of the best possible order of the bound with respect to the dimension, let us notice that, from Jensen's inequality,

$$\widetilde{\gamma}_+ \geq \frac{1}{p+1}\left(\frac{2\big[\log(\epsilon^{-1}) + 0.73\,d\big]}{3(\kappa-1)n}\right)^{p/2}\big[(1+\widehat{\delta})d - 1\big]^{p+1}\big[1 + (n\epsilon)^{-1/2}\big].$$

Therefore, if we take for example $\epsilon = 1/n$, the best we can hope for is to get a $\widetilde{\gamma}_+$ of order $d^{3p/2+1}/n^{p/2} = d^4/n$ when $p=2$. The power of d in this second

order term is clearly not optimal, due to the rather crude inequality (9) used to get this proposition. It would have been more satisfactory to get an upper bound $\widetilde{\gamma}_+$ of the same order as γ_-, that is of order d/n, at least when $p = 2$. Anyhow, the idea here is that, since we are dealing with a second order term in $1/n$, we can privilege the simplicity of the proof over the sharpness of the result.

Also, one may think that imposing that $\mathbb{E}\big(\|G^{-1/2}X\|^{4p+4}\big) < \infty$ for some $p > 1$ that is necessary to get $\widetilde{\gamma}_+ < \infty$ is asking for a pretty high moment condition. One can get a rate depending on a lower moment assumption using the following variant of Bienaymé Chebishev's inequality.

LEMMA 2.5 *Let $q \in [1, 2]$ be some exponent. Consider the constant*

$$C_q = \frac{q^{q-1}}{2(q-1)^{q-1}(1-q/2)^{(2-q)/q}} \leq 1.4,$$

where by convention $C_1 = \lim_{q\to 1+} C_q = 1$ and $C_2 = \lim_{q\to 2-} C_q = 1$. Let $W_1, \ldots, W_n$ be n independent copies of a non-negative real valued random variable W. With probability at least $1 - 2\epsilon$,

$$\frac{1}{n}\sum_{i=1}^{n} W_i \leq \mathbb{E}(W) + \frac{C_q \mathbb{E}(W^q)^{1/q}}{\epsilon^{1/q} n^{1-1/q}}.$$

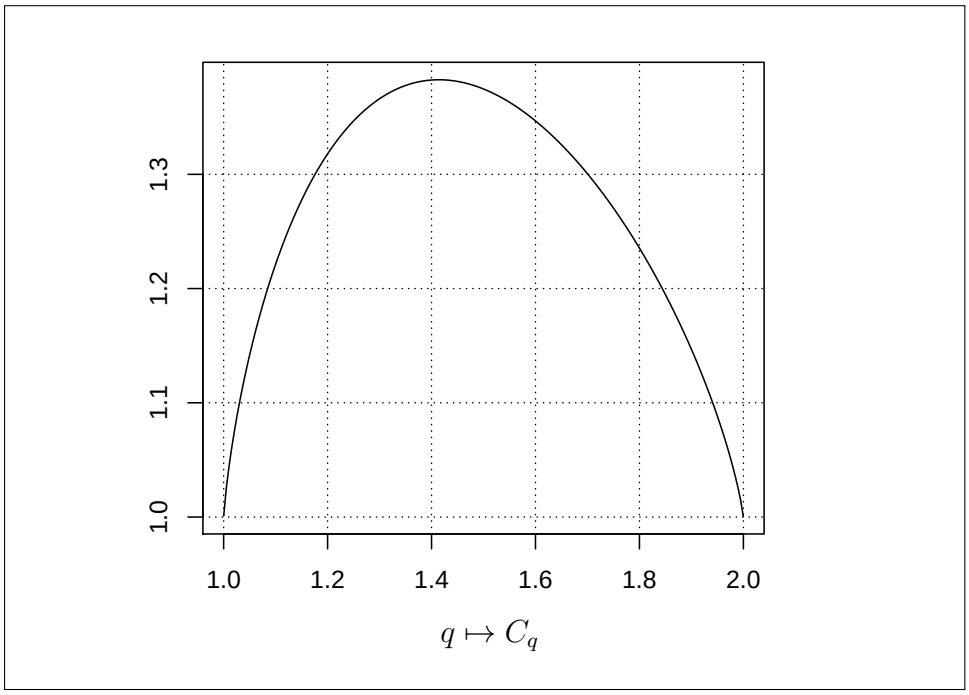

Using this lemma, we obtain the following variant of Proposition 2.4.

PROPOSITION 2.6 *Consider some exponents $p \in [1,2]$ and $q \in [1,2]$ and the bound*

$$\widetilde{\gamma}_+ = \frac{1}{p+1}\left(\frac{2\big[\log(\epsilon^{-1}+0.73d\big]}{(\kappa-1)n}\right)^{p/2}\big(1+\widehat{\delta}\big)^{p+1} \times \left[\mathbb{E}\big(\|G^{-1/2}X\|^{2(p+1)} + \frac{C_q\mathbb{E}\Big(\|G^{-1/2}X\|^{2q(p+1)}\Big)^{1/q}}{\epsilon^{1/q}n^{1-1/q}}\right].$$

Under the hypotheses of Proposition 1.1 on page 7, with probability at least $1-4\epsilon$, for any $\theta \in \mathbb{R}^d$,

$$-\underbrace{\frac{\widehat{\delta}+\gamma_-}{1+\widehat{\delta}}}_{=\mathbf{O}(\widehat{\delta})} \le \frac{\overline{N}(\theta)}{N(\theta)} - 1 \le \frac{1}{(1-\widehat{\delta})(1-\widetilde{\gamma}_+)_+} - 1 \le \underbrace{\frac{\widehat{\delta}+\widetilde{\gamma}_+}{(1-\widehat{\delta})(1-\widetilde{\gamma}_+)_+}}_{\substack{=\mathbf{O}_{n\to\infty}(\widehat{\delta}) \\ \textit{when } p>1}},$$

where it is useful to remember that

$$\widehat{\delta} = \mathbf{O}\left(\sqrt{\frac{\kappa\big[d+\log(\epsilon^{-1})\big]}{n}}\right).$$

We have done two things in this section about the empirical Gram matrix estimate. The first was to make a direct comparison between $\overline{N}$ and N. The second was to introduce non-random bounds in terms of the quantities γ_-, $\widehat{\gamma}_+$ and $\widetilde{\gamma}_+$.

Remark that γ_-, controlling the accuracy of $\overline{N}(\theta)$ as an upper bound for $N(\theta)$, is always of order $1/n$, meaning that the empirical quadratic form $\overline{N}(\theta)$ always provides an upper bound of the same quality as the robust estimate $\widehat{N}(\theta)$. This comes from the fact that $\langle\theta, X\rangle^2$ being a non-negative random variable cannot have a long tail on the left-hand side, but only on the right-hand side, so that its empirical mean can be too large, but not too small.

On the other hand, $\delta_+(\theta)$, controlling the accuracy of $\overline{N}(\theta)$ as a lower bound for $N(\theta)$, is not necessarily a second order term in all circumstances. When this is not case, the fact that the empirical Gram estimate may become less accurate than the robust estimate $\widehat{N}(\theta)$ is not ruled out, and we will see on simulations that this does happen in practice.

3. Estimation of the covariance matrix

To estimate the covariance matrix

$$\Sigma = \mathbb{E}\big[\big(X - \mathbb{E}(X)\big)\big(X - \mathbb{E}(X)\big)^\top\big],$$

one can remark that

$$\theta^\top \Sigma \theta = \inf_{\gamma \in \mathbb{R}} \mathbb{E}\big[\big(\langle \theta, X\rangle - \gamma\big)^2\big]$$

and consider the quadratic form

$$N(\theta, \gamma) = \mathbb{E}\big[\big[\langle \theta, X\rangle - \gamma\big)^2\big]$$

that corresponds to the Gram matrix of the extended variable $(X, -1) \in \mathbb{R}^{d+1}$.

Proposition 3.1 *Consider an estimator $\widehat{N}(\theta, \gamma)$ of $N(\theta, \gamma)$. Let δ and ϵ be two positive constants. Assume that with probability at least $1 - \epsilon$, for any $(\theta, \gamma) \in \mathbb{R}^{d+1}$*

$$\left| \frac{N(\theta, \gamma)}{\widehat{N}(\theta, \gamma)} - 1 \right| \leq \delta.$$

Define

$$\widetilde{N}(\theta) = \inf_{\gamma \in \mathbb{R}} \widehat{N}(\theta, \gamma)$$

With probability at least $1 - \epsilon$, for any $\theta \in \mathbb{R}^d$,

$$\left| \frac{\theta^\top \Sigma \theta}{\widetilde{N}(\theta)} - 1 \right| \leq \delta.$$

Proof. With probability at least $1 - \epsilon$,

$$\begin{aligned}(1-\delta)\widetilde{N}(\theta) = \inf_{\gamma \in \mathbb{R}} (1-\delta)\widehat{N}(\theta, \gamma) &\leq \inf_{\gamma \in \mathbb{R}} N(\theta, \gamma) \\ &= \theta^\top \Sigma \theta \leq \inf_{\gamma} (1+\delta)\widehat{N}(\theta, \gamma) = (1+\delta)\widetilde{N}(\theta).\end{aligned}$$

□

We can then use the previous sections to describe more precisely estimators $\widehat{N}(\theta, \gamma)$ that can be used under suitable conditions. To establish those conditions, the following lemma will be helpful.

LEMMA 3.2 *Let us put*

$$\kappa = \sup\Big\{\mathbb{E}\big(\langle\theta, X - \mathbb{E}(X)\rangle^4\big) \,:\, \theta \in \mathbb{R}^d,\ \mathbb{E}\big(\langle\theta, X - \mathbb{E}(X)\rangle^2\big) \leq 1\Big\}$$

and $$\kappa' = \sup\Big\{\mathbb{E}\big[\big(\langle\theta, X\rangle - \xi\big)^4\big] \,:\, \theta \in \mathbb{R}^d, \xi \in \mathbb{R},\ \mathbb{E}\big[\big(\langle\theta, X\rangle - \xi\big)^2\big] \leq 1\Big\}.$$

These two kurtosis coefficients are related by the inequality

$$\kappa' \leq (\sqrt{\kappa} + 1)^2.$$

PROOF. Using successively the triangular inequality in $\mathbb{L}^4(\mathbb{P})$, the definition of κ and the Cauchy-Schwarz inequality in $\mathbb{R}^2$,

$$\begin{aligned}
\mathbb{E}\big[\big(\langle\theta, X\rangle - \xi\big)^4\big]^{1/2} &\leq \Big(\mathbb{E}\big(\langle\theta, X - \mathbb{E}(X)\rangle^4\big)^{1/4} + |\langle\theta, \mathbb{E}(X)\rangle - \xi|\Big)^2 \\
&\leq \Big(\kappa^{1/4}\mathbb{E}\big(\langle\theta, X - \mathbb{E}(X)\rangle^2\big)^{1/2} + |\langle\theta, \mathbb{E}(X)\rangle - \xi|\Big)^2 \\
&\leq (\kappa^{1/2} + 1)\Big[\mathbb{E}\big(\langle\theta, X - \mathbb{E}(X)\rangle^2\big) + \big(\langle\theta, \mathbb{E}(X)\rangle - \xi\big)^2\Big] \\
&= (\kappa^{1/2} + 1)\mathbb{E}\big[\big(\langle\theta, X\rangle - \xi\big)^2\big].
\end{aligned}$$

□

PROPOSITION 3.3 *Let $(X_1, \ldots, X_n)$ be n independent copies of a vector valued random variable $X \in \mathbb{R}^d$. Assume that for any $\theta \in \mathbb{R}^d$,*

$$\mathbb{E}\big(\langle\theta, X - \mathbb{E}(X)\rangle^4\big) \leq \kappa\, \mathbb{E}\big(\langle\theta, X - \mathbb{E}(X)\rangle^2\big)^2,$$

for a known constant $\kappa \in \mathbb{R}_+$. Let

$$\widehat{N}(\theta) = \inf_{\xi\in\mathbb{R}} \inf\left\{\rho \in \mathbb{R}_+^* \,,\, \sum_{i=1}^n \psi\left\{\lambda\left[\frac{\big(\langle\theta, X_i\rangle - \xi\big)^2}{\rho} - 1\right]\right\} \leq 0\right\},$$

where

$$\lambda = \sqrt{\frac{2}{(\kappa + 2\sqrt{\kappa})n}\big[\log(\epsilon^{-1}) + 0.73(d+1)\big]}.$$

For any confidence parameter $\epsilon > 0$, and any sample size n such that

$$\begin{aligned}
n > \Bigg[&20\,(\kappa^{1/2} + 1)(d+1)^{1/2} \\
&+ \left(\frac{5}{2} + \frac{1}{2(\kappa + 2\kappa^{1/2})}\right)\sqrt{2\,(\kappa + 2\kappa^{1/2})\big[\log(\epsilon^{-1}) + 0.73(d+1)\big]}\Bigg]^2, \quad (10)
\end{aligned}$$

with probability at least $1-2\epsilon$*, for any* $\theta \in \mathbb{R}^d$*,*

$$\left|\frac{\mathbb{E}\big(\langle\theta, X-\mathbb{E}(X)\rangle^2\big)}{\widehat{N}(\theta)}-1\right| \leq \frac{\mu}{1-2\mu} = \mathbf{O}\left(\sqrt{\frac{\kappa\big[d+\log(\epsilon^{-1})\big]}{n}}\right),$$

where

$$\mu = \sqrt{\frac{2(\kappa+2\kappa^{1/2})}{n}\big[\log(\epsilon^{-1})+0.73(d+1)\big]} + 6.81\,(\kappa^{1/2}+1)\sqrt{\frac{2(d+1)}{n}},$$

assuming by convention that $0/0=1$ *and* $z/0=+\infty$ *for any* $z>0$*.*

The proof of this proposition is given in the appendix.

Consider now the empirical covariance estimate

$$\overline{\Sigma} = \frac{1}{2n^2}\sum_{i=1}^n\sum_{j=1}^n (X_i-X_j)(X_i-X_j)^\top.$$

(We choose a biased normalization by n^2 instead of $n(n-1)$, because, as we will see, we can prove a simpler non-asymptotic result for it.)

Remark that

$$\theta^\top\overline{\Sigma}\,\theta = \inf_{\xi\in\mathbb{R}}\frac{1}{n}\sum_{i=1}^n\big(\langle\theta,X_i\rangle-\xi\big)^2 = \frac{1}{n}\sum_{i=1}^n\Big\langle\theta, X_i-\frac{1}{n}\sum_{j=1}^n X_j\Big\rangle^2.$$

In order to use Proposition 2.2 on page 11, we need the following lemma.

Lemma 3.4 *Almost surely, for any* $(\theta,\xi)\in\mathbb{R}^{d+1}$ *such that*

$$\mathbb{E}\big[\big(\langle\theta,X\rangle-\xi\big)^2\big] \leq 1,$$

$$\langle\theta,X_i\rangle-\xi \leq \big(\|\Sigma^{-1/2}\big(X_i-\mathbb{E}(X)\big)\|^2+1\big)^{1/2}.$$

This lemma is proved in appendix.

We are now ready to apply Proposition 2.2 on page 11. Define

$$\begin{aligned}
R &= \max_{i=1,\dots,n}\big(\big\|\Sigma^{-1/2}\big(X_i-\mathbb{E}(X)\big)\big\|^2+1\big)^{1/2},\\
\mu &= \sqrt{\frac{2(\kappa+2\kappa^{1/2})}{n}\big[\log(\epsilon^{-1})+0.73\,(d+1)\big]} + 6.81\,(\kappa^{1/2}+1)\sqrt{\frac{2(d+1)}{n}},\\
\widehat{\delta} &= \frac{\mu}{1-2\mu},\\
\gamma_- &= \frac{2\big[\log(\epsilon^{-1})+0.73\,(d+1)\big]}{3(\kappa+2\kappa^{1/2})\,n},\\
\gamma_+ &= \gamma_-\,R^4\big(1+\widehat{\delta}\big)^2.
\end{aligned}$$

PROPOSITION 3.5 *Assume that for some known constant* $\kappa \in \mathbb{R}_+$,

$$\sup\Big\{\mathbb{E}\big(\langle\theta, X - \mathbb{E}(X)\rangle^4\big) \,:\, \theta \in \mathbb{R}^d, \mathbb{E}\big(\langle\theta, X - \mathbb{E}(X)\rangle^2\big) \leq 1\Big\} \leq \kappa < \infty.$$

Consider any confidence parameter $\epsilon > 0$ *and any sample size* n *satisfying*

$$n > \left[20(\kappa^{1/2}+1)\sqrt{d+1} + \left(\frac{5}{2} + \frac{1}{2(\kappa+2\kappa^{1/2})}\right)\sqrt{2(\kappa+2\kappa^{1/2})\big[\log(\epsilon^{-1}) + 0.73\,(d+1)\big]}\right]^2.$$

With probability at least $1-2\epsilon$, *for any* $\theta \in \mathbb{R}^d$,

$$-\underbrace{\frac{\widehat{\delta}+\gamma_-}{1+\widehat{\delta}}}_{=\mathbf{O}(\widehat{\delta})} \leq \frac{\theta^\top\overline{\Sigma}\theta}{\theta^\top\Sigma\theta} - 1 \leq \frac{\widehat{\delta}+\gamma_+}{(1-\widehat{\delta})(1-\gamma_+)_+},$$

where $\widehat{\delta} = \mathbf{O}\left(\sqrt{\dfrac{\kappa\big[d+\log(\epsilon^{-1})\big]}{n}}\right)$.

PROOF. Let us introduce

$$\overline{N}(\theta,\xi) = \frac{1}{n}\sum_{i=1}^n \big(\langle\theta, X_i\rangle - \xi\big)^2.$$

From proposition 2.2 on page 11, we see that on some event Ω' of probability at least $1-2\epsilon$, for any $(\theta,\xi) \in \mathbb{R}^{d+1}$,

$$B_- \stackrel{\text{def}}{=} -\frac{\widehat{\delta}+\gamma_-}{1+\widehat{\delta}} \leq \frac{\overline{N}(\theta,\xi)}{N(\theta,\xi)} - 1 \leq B_+ \stackrel{\text{def}}{=} \frac{\widehat{\delta}+\gamma_+}{(1-\widehat{\delta})(1-\gamma_+)_+}.$$

Remark that on Ω',

$$\theta^\top\overline{\Sigma}\theta = \inf_{\xi\in\mathbb{R}} \overline{N}(\theta,\xi) \leq (1+B_+)\inf_{\xi\in\mathbb{R}} N(\theta,\xi) = (1+B_+)\,\theta^\top\Sigma\theta,$$
$$\theta^\top\overline{\Sigma}\theta \geq (1-B_-)\inf_{\xi\in\mathbb{R}} N(\theta,\xi) = (1-B_-)\,\theta^\top\Sigma\theta.$$

□

It is easy from there to give non-asymptotic bounds for the empirical term R. The discussion is similar to the one following equation (7) on page 12, with G replaced by Σ and X replaced by $X - \mathbb{E}(X)$, so that we will not repeat it. One can also obtain results similar to Propositions 2.4 on page 15 and Proposition 2.6 on page 17.

4. Least squares regression

In this section, we use our results on the estimation of the Gram matrix in two ways. First, we derive new robust estimators for least squares regression with a random design. Second, we obtain new results for the ordinary least squares estimator, including its exact asymptotic rate of convergence in expectation under quite weak assumptions. In particular, it turns out that this rate is C/n, with a constant C different from $d\sigma^2$, where σ^2 is the variance of the noise, in the case when the noise is correlated to the design.

4.1. From Gram matrix to least squares estimates. Let us first make a connection between Gram matrix estimates and least squares estimates, that will serve to study a new robust least squares estimate as well as the empirical risk minimizer. Consider n independent copies (X_i, Y_i), $i = 1, \dots, n$, of the couple of random variables $(X, Y) \in \mathbb{R}^d \times \mathbb{R}$ and the question of minimizing the quadratic risk

$$R(\theta) = \mathbb{E}\big[\big(Y - \langle \theta, X\rangle\big)^2\big]$$

in $\theta \in \mathbb{R}^d$. Introduce the homogeneous quadratic form

$$N(\theta, \xi) = \mathbb{E}\big[\big(\xi Y - \langle \theta, X\rangle\big)^2\big], \quad \theta \in \mathbb{R}^d, \xi \in \mathbb{R},$$

and assume that $\breve{N}(\theta, \xi)$ is a quadratic estimator of $N(\theta, \xi)$, based on the sample (X_i, Y_i), $i = 1, \dots, n$, such that, for some positive constant δ, with probability $1 - \epsilon$, for any $(\theta, \xi) \in \mathbb{R}^{d+1}$,

$$\left| \frac{\breve{N}(\theta, \xi)}{N(\theta, \xi)} - 1 \right| \leq \delta. \tag{11}$$

Hypotheses under which such an estimator exists are stated in Proposition 1.2 on page 8 for a robust estimator and in Proposition 2.3 on page 14 and 2.4 on page 15 for the empirical Gram matrix estimator. We will come back to this in more detail afterwards.

Proposition 4.1 *Under the above hypotheses, any estimator*

$$\widehat{\theta} \in \arg\min_{\theta \in \mathbb{R}^d} \breve{N}(\theta, 1)$$

is such that with probability at least $1 - \epsilon$,

$$R(\widehat{\theta}) - \inf_{\theta \in \mathbb{R}^d} R(\theta) \leq \frac{\delta^2}{(1-\delta)(1-\delta^2)} \breve{N}(\widehat{\theta}, 1).$$

Corollary 4.2 *With probability at least* $1-\epsilon$*, for any* $\theta_* \in \arg\min_{\theta\in\mathbb{R}^d} R(\theta)$*,*

$$\begin{aligned}\frac{\breve{N}(\widehat{\theta}-\theta_*,0)}{1+\delta} \le N(\widehat{\theta}-\theta_*,0) = R(\widehat{\theta}) - R(\theta_*) &\le \frac{\delta^2}{(1-\delta)(1-\delta^2)}\breve{N}(\widehat{\theta},1)\\ \le \frac{\delta^2}{(1-\delta)(1-\delta^2)}\breve{N}(\theta_*,1) &\le \frac{\delta^2}{(1-\delta)^2}R(\theta_*) \le \frac{\delta^2}{(1-\delta)^2}R(0)\\ &= \frac{\delta^2}{(1-\delta)^2}\mathbb{E}\big(Y^2\big).\end{aligned}$$

This gives an observable non-asymptotic confidence region for θ_**, defined by the equation*

$$\breve{N}(\widehat{\theta}-\theta_*,0) \le \frac{\delta^2}{(1-\delta)^2}\breve{N}(\widehat{\theta},1).$$

This also provides a simple non-asymptotic bound for the excess risk

$$R(\widehat{\theta}) - R(\theta_*) \le \frac{\delta^2}{(1-\delta)^2}R(\theta_*) \le \frac{\delta^2}{(1-\delta)^2}\mathbb{E}\big(Y^2\big).$$

As a consequence, in the case when condition (11) *on page 22 is satisfied and* $Y = f(X) + W$ *where* W *is independent of* X*, centered and* $\mathbb{E}(W^2) \le \sigma^2$*, with probability at least* $1-\epsilon$*,*

$$R(\widehat{\theta}) - R(\theta_*) \le \frac{\delta^2}{(1-\delta)^2}\Big\{\sigma^2 + \mathbb{E}\big[f(X)^2\big]\Big\}.$$

Proof. We take advantage of the fact that not only the derivatives of $\theta \mapsto R(\theta)$ and $\theta \mapsto \breve{N}(\theta,1)$ vanish at θ_* and $\widehat{\theta}$, respectively, but also their symmetric finite differences $R(\theta+\theta') - R(\theta-\theta')$ and $\breve{N}\big(\theta+\theta',1\big) - \breve{N}\big(\theta-\theta',1\big)$.

More precisely, using the Cauchy-Schwarz inequality (as well as its case of equality), the fact that $\mathbb{E}\big[\big(\langle\theta_*,X\rangle - Y\big)\langle\theta',X\rangle\big] = 0$ for any $\theta' \in \mathbb{R}^d$, and equation (11), we obtain the following chain of inequalities, that holds with probability at least $1-\epsilon$, for any (possibly random) positive number a:

$$\begin{aligned}R(\widehat{\theta}) - R(\theta_*) &= \mathbb{E}\big(\langle\widehat{\theta}-\theta_*,X\rangle^2\big)\\ &= \sup\Big\{a^{-1}\mathbb{E}\big(\langle\widehat{\theta}-\theta_*,X\rangle\langle\theta',X\rangle\big) \,:\, \theta' \in \mathbb{R}^d, \mathbb{E}\big(\langle\theta',X\rangle^2\big) \le a^2\Big\}^2\\ &= \sup\Big\{a^{-1}\mathbb{E}\big[\big(\langle\widehat{\theta},X\rangle - Y\big)\langle\theta',X\rangle\big] \,|\, \theta' \in \mathbb{R}^d, \mathbb{E}\big(\langle\theta',X\rangle^2\big) \le a^2\Big\}^2\\ &= \sup\Big\{\frac{1}{4a}\Big[\mathbb{E}\Big(\big(\langle\widehat{\theta}+\theta',X\rangle - Y\big)^2\Big) - \mathbb{E}\Big(\big(\langle\widehat{\theta}-\theta',X\rangle - Y\big)^2\Big)\Big]\\ &\qquad : \theta' \in \mathbb{R}^d, \mathbb{E}\big(\langle\theta',X\rangle^2\big) \le a^2\Big\}^2\end{aligned}$$

$$\leq \sup\Bigg\{ \frac{1}{4a}\Bigg[\frac{\breve{N}\big(\widehat{\theta}+\theta',1\big)}{1-\delta} - \frac{\breve{N}\big(\widehat{\theta}-\theta',1\big)}{1+\delta}\Bigg] : \theta' \in \mathbb{R}^d, \breve{N}(\theta',0) \leq a^2(1+\delta)\Bigg\}^2$$

$$= \sup\Bigg\{ \frac{\delta}{2a(1-\delta^2)}\Big[\breve{N}\big(\widehat{\theta},1\big) + \breve{N}\big(\theta',0\big)\Big] : \theta' \in \mathbb{R}^d, \breve{N}\big(\theta',0\big) \leq a^2(1+\delta)\Bigg\}^2$$

$$= \frac{\delta^2}{4(1-\delta^2)^2}\Bigg[\frac{\breve{N}\big(\widehat{\theta},1\big)}{a} + a(1+\delta)\Bigg]^2$$

Taking the optimal value $a = \sqrt{\dfrac{\breve{N}\big(\widehat{\theta},1\big)}{1+\delta}}$, we obtain as desired that with probability at least $1-\epsilon$,

$$R(\widehat{\theta}) - R(\theta_*) \leq \frac{\delta^2}{(1-\delta)(1-\delta^2)}\breve{N}\big(\widehat{\theta},1\big).$$

The statements made in the corollary are obvious consequences of this inequality and the definitions. Note that this line of proof would also work (with the necessary modifications) in the case when (X_i, Y_i) are independent couples of random variables that are not necessarily identically distributed. □

4.2. A ROBUST LEAST SQUARES ESTIMATOR. To draw the consequences of Proposition 4.1, we have to explain when the required hypothesis, expressed by equation (11) on page 22, is satisfied.

For this, we will apply Proposition 1.2 on page 8, assuming that

$$\sup\Big\{\mathbb{E}\Big[\big(\xi Y - \langle\theta, X\rangle\big)^4\Big] : \theta \in \mathbb{R}^d, \xi \in \mathbb{R}, \mathbb{E}\Big[\big(\xi Y - \langle\theta, X\rangle\big)^2\Big] \leq 1\Big\} \leq \kappa < \infty. \quad (12)$$

When this is satisfied, there is a robust estimator satisfying condition (11) with $\delta = 2\mu/(1-4\mu)$, where μ is as in Proposition 1.1 on page 7 with d replaced by $d+1$, and is therefore of order $\mathbf{O}\Big(\sqrt{\kappa\big[d+\log(\epsilon^{-1})\big]/n}\Big)$.

Condition (12) is not very explicit, since it bears on the joint distribution of X and Y. It may be more instructive to replace it by separate kurtosis assumptions bearing on $Y - \langle\theta_*, X\rangle$ and on $\langle\theta, X\rangle$, where $\theta_* \in \arg\min_{\theta\in\mathbb{R}^d} \mathbb{E}\big[\big(Y - \langle\theta, X\rangle\big)^2\big]$. This is what the following lemma does.

Lemma 4.3 *Let us define*

$$\kappa = \sup\Big\{\mathbb{E}\big[\big(\xi Y - \langle\theta, X\rangle\big)^4\big] \,:\, \theta\in\mathbb{R}^d, \xi\in\mathbb{R},\ \mathbb{E}\big[\big(\xi Y - \langle\theta, X\rangle\big)^2\big] \le 1\Big\},$$

$$\kappa_1 = \sup\Big\{\mathbb{E}\big(\langle\theta, X\rangle^4\big) \,:\, \theta\in\mathbb{R}^d,\ \mathbb{E}\big(\langle\theta, X\rangle^2\big) \le 1\Big\},$$

$$\kappa_2 = \begin{cases} \dfrac{\mathbb{E}\big[\big(Y - \langle\theta_*, X\rangle\big)^4\big]}{\mathbb{E}\big[\big(Y - \langle\theta_*, X\rangle\big)^2\big]^2}, & \mathbb{E}\big[\big(Y - \langle\theta_*, X\rangle\big)^2\big] > 0, \\ 0, & \text{otherwise.} \end{cases}$$

Those three kurtosis coefficients are linked together by the relation

$$\sqrt{\kappa} \le \sqrt{\kappa_1} + \sqrt{\kappa_2}.$$

Proof. Using the triangular inequality in $\mathbb{L}^4$, followed by the definitions of κ_1 and κ_2 and the Cauchy-Schwarz inequality in $\mathbb{R}^2$, we see that

$$\begin{aligned}
\mathbb{E}\big[\big(\xi Y - \langle\theta, X\rangle\big)^4\big]^{1/2} &= \mathbb{E}\Big\{\big[\xi\big(Y - \langle\theta_*, X\rangle\big) + \langle\xi\theta_* - \theta, X\rangle\big]^4\Big\}^{1/2} \\
&\le \Big\{|\xi|\mathbb{E}\big[\big(Y - \langle\theta_*, X\rangle\big)^4\big]^{1/4} + \mathbb{E}\big(\langle\xi\theta_* - \theta, X\rangle^4\big)^{1/4}\Big\}^2 \\
&\le \Big\{|\xi|\kappa_2^{1/4}\mathbb{E}\big[\big(Y - \langle\theta_*, X\rangle\big)^2\big]^{1/2} + \kappa_1^{1/4}\mathbb{E}\big(\langle\xi\theta_* - \theta, X\rangle^2\big)^{1/2}\Big\}^2 \\
&\le (\kappa_1^{1/2} + \kappa_2^{1/2})\Big\{\xi^2\mathbb{E}\big[\big(Y - \langle\theta_*, X\rangle\big)^2\big] + \mathbb{E}\big(\langle\xi\theta_* - \theta, X\rangle^2\big)\Big\} \\
&= (\kappa_1^{1/2} + \kappa_2^{1/2})\mathbb{E}\big[\big(\xi Y - \langle\theta, X\rangle\big)^2\big].
\end{aligned}$$

□

As a consequence of this lemma and of Proposition 1.2 on page 8

Proposition 4.4 *Consider the bound*

$$\mu = \sqrt{\frac{2\big[\big(\sqrt{\kappa_1} + \sqrt{\kappa_2}\big)^2 - 1\big]\big[\log(\epsilon^{-1}) + 0.73\,(d+1)\big]}{n}} + 6.81\,\big(\sqrt{\kappa_1} + \sqrt{\kappa_2}\big)\sqrt{\frac{2(d+1)}{n}}.$$

There exists a robust quadratic estimator $\breve{N}(\theta, \xi)$ such that for any estimator

$$\widehat{\theta} \in \arg\min_{\theta\in\mathbb{R}^d} \breve{N}(\theta, 1),$$

for any $\epsilon > 0$ and $n \in \mathbb{N}$ satisfying

$$n > \left[20\big(\sqrt{\kappa_1}+\sqrt{\kappa_2}\big)\sqrt{d+1} + \left(\frac{5}{2} + \frac{1}{2\big[\big(\sqrt{\kappa_1}+\sqrt{\kappa_2}\big)^2 - 1\big]}\right) \times \left(2\big[\big(\sqrt{\kappa_1}+\sqrt{\kappa_2}\big)^2 - 1\big]\big[\log(\epsilon^{-1}) + 0.73\,(d+1)\big]\right)^{1/2}\right]^2 = \mathbf{O}\Big((\kappa_1+\kappa_2)\big[d+\log(\epsilon^{-1})\big]\Big),$$

with probability at least $1-2\epsilon$,

$$R(\widehat{\theta}) - R(\theta_*) \leq \frac{\delta^2}{(1-\delta)^2} R(\theta_*),$$

where

$$\delta^2 = \left(\frac{2\mu}{1-4\mu}\right)^2 = \mathbf{O}\left(\frac{(\kappa_1+\kappa_2)\big[d+\log(\epsilon^{-1})\big]}{n}\right).$$

4.3. Generalization bounds for the empirical risk minimizer. Let us now examine the conditions under which equation (11) on page 22 is satisfied by the empirical Gram matrix estimator

$$\breve{N}(\theta,\xi) = \overline{N}(\theta,\xi) = \frac{1}{n}\sum_{i=1}^n \big(\xi Y_i - \langle \theta, X_i\rangle\big)^2.$$

To apply Proposition 2.2 on page 11, we have to bound

$$R = \max_{i=1,\dots,n} \sup\Big\{\xi Y_i - \langle\theta, X_i\rangle \,:\, (\theta,\xi)\in\mathbb{R}^{d+1}, \mathbb{E}\big[\big(\xi Y - \langle\theta, X\rangle\big)^2\big] \leq 1\Big\}. \quad (13)$$

Lemma 4.5 *The above defined quantity satisfies almost surely*

$$R^2 \leq \max_{i=1,\dots,n}\left\{\frac{\big(Y_i - \langle\theta_*, X_i\rangle\big)^2}{\mathbb{E}\big[\big(Y-\langle\theta_*,X\rangle\big)^2\big]} + \big\|G^{-1/2}X_i\big\|^2\right\}, \leq \max_{i=1,\dots,n}\frac{\big(Y_i - \langle\theta_*, X_i\rangle\big)^2}{\mathbb{E}\big[\big(Y-\langle\theta_*,X\rangle\big)^2\big]} + \max_{i=1,\dots,n}\big\|G^{-1/2}X_i\big\|^2,$$

with the convention that $0/0 = 0$.

Proof. Remark that for any positive constants a and b,

$$\begin{aligned}\big(\xi Y_i - \langle\theta, X_i\rangle\big)^2 &= \big[\xi\big(Y_i - \langle\theta_*, X_i\rangle\big)a^{-1}a + \langle\xi\theta_* - \theta, X_i\rangle b^{-1}b\big]^2 \\ &\leq \big[\xi^2\big(Y_i - \langle\theta_*, X_i\rangle\big)^2 a^{-2} + \langle\xi\theta_* - \theta, X_i\rangle^2 b^{-2}\big]\big(a^2 + b^2\big).\end{aligned}$$

Now take $a^2 = \xi^2\mathbb{E}\big[\big(Y - \langle\theta_*, X\rangle\big)^2\big]$ and $b^2 = \mathbb{E}\big(\langle\xi\theta_* - \theta, X\rangle^2\big)$, and remark that

$$a^2 + b^2 = \mathbb{E}\big[\big(\xi Y - \langle\theta, X\rangle\big)^2\big].$$

Notice also that in the case when $b^2 > 0$,

$$\langle\xi\theta_* - \theta, X_i\rangle^2 b^{-2} \leq \|G^{1/2}(\xi\theta_* - \theta)\|^2 \|G^{-1/2}X_i\|^2 b^{-2} = \|G^{-1/2}X_i\|^2.$$

Consequently, when $a^2 > 0$ and $b^2 > 0$,

$$\begin{aligned}\big(\xi Y_i - \langle\theta, X_i\rangle\big)^2 \leq \Bigg(&\frac{\big(Y_i - \langle\theta_*, X_i\rangle\big)^2}{\mathbb{E}\big[\big(Y - \langle\theta_*, X\rangle\big)^2\big]} \\ &+ \|G^{-1/2}X_i\|^2\Bigg)\mathbb{E}\big[\big(\xi Y - \langle\theta, X\rangle\big)^2\big]. \quad (14)\end{aligned}$$

In the case when $a^2 = 0$, $Y_i - \langle\theta_*, X_i\rangle = 0$ almost surely, and in the case when $b = 0$, $\langle\xi\theta_* - \theta, X_i\rangle = 0$ almost surely, so that in those cases, equation (14) is still satisfied, and the desired result is an easy consequence of this inequality. □

In view of the above lemma, suitable hypotheses to obtain a bound for R are that $\mathbb{E}\big[\big(Y - \langle\theta_*, X\rangle\big)^2\big] > 0$ and for positive constants $p, q \in]0, 1]$, α_i and η_i, $i = 1, 2$,

$$\mathbb{E}\Big\{\exp\Big[\frac{\alpha_1}{2}\Big(\|G^{-1/2}X\|^{2p} - d^p - \eta_1\Big)\Big]\Big\} \leq 1 \quad (15)$$

$$\text{and } \mathbb{E}\Bigg\{\exp\Bigg[\frac{\alpha_2}{2}\Bigg(\frac{(Y - \langle\theta_*, X\rangle)^{2q}}{\mathbb{E}\big[(Y - \langle\theta_*, X\rangle)^2\big]^q} - 1 - \eta_2\Bigg)\Bigg]\Bigg\} \leq 1. \quad (16)$$

Under those hypotheses, with probability at least $1 - 2\epsilon$,

$$R^2 \leq \Bigg(d^p + \eta_1 + \frac{2}{\alpha_1}\log\big(n/\epsilon\big)\Bigg)^{1/p} + \Bigg(1 + \eta_2 + \frac{2}{\alpha_2}\log\big(n/\epsilon\big)\Bigg)^{1/q}.$$

In the case when $\mathbb{E}\big[\big(Y - \langle\theta_*, X\rangle\big)^2\big] = 0$, it is easy to see that when (15) is satisfied, with probability at least $1 - \epsilon$

$$R^2 \leq \Bigg(d^p + \eta_1 + \frac{2}{\alpha_1}\log(n/\epsilon)\Bigg)^{1/p}.$$

Applying Proposition 2.2 on page 11 and Proposition 4.1 on page 22 and its corollary, we obtain

Proposition 4.6 *Assume that the hypotheses expressed by equations* (15) *and* (16) *are satisfied. Consider*

$$\begin{aligned}
\mu &= \left(\frac{2\left[\left(\sqrt{\kappa_1}+\sqrt{\kappa_2}\right)^2-1\right]}{n}\left[\log(\epsilon^{-1})+0.73\,(d+1)\right]\right)^{1/2} \\
&\quad + 6.81\left(\sqrt{\kappa_1}+\sqrt{\kappa_2}\right)\sqrt{\frac{2(d+1)}{n}}, \\
\widehat{\delta} &= \frac{\mu}{1-2\mu}, \\
\gamma_- &= \frac{2\left[\log(\epsilon^{-1})+0.73\,(d+1)\right]}{3\left[(\sqrt{\kappa_1}+\sqrt{\kappa_2})^2-1\right]n}, \\
\gamma_+ &= \gamma_1 R^4\left(1+\widehat{\delta}\right)^2, \\
\widehat{\gamma}_+ &= \frac{2\left[\log(\epsilon^{-1})+0.73\,(d+1)\right](1+\widehat{\delta})^2}{3\left[(\sqrt{\kappa_1}+\sqrt{\kappa_2})^2-1\right]n} \\
&\quad \times\left[\left(d^p+\eta_1+2\alpha_1^{-1}\log(n/\epsilon)\right)^{2/p}+\left(1+\eta_2+2\alpha_2^{-1}\log(n/\epsilon)\right)^{2/q}\right]^2,
\end{aligned}$$

where κ_1 *and* κ_2 *are defined in Lemma 4.3 on page 25 and* R *is defined in* (13) *on page 26. Consider any confidence parameter* $\epsilon>0$ *and any sample size* n *such that*

$$\begin{aligned}
n &> \left[20\left(\sqrt{\kappa_1}+\sqrt{\kappa_2}\right)\sqrt{d+1}+\left(\frac{5}{2}+\frac{1}{2\left[\left(\sqrt{\kappa_1}+\sqrt{\kappa_2}\right)^2-1\right]}\right)\right. \\
&\quad \left.\times\left(2\left[\left(\sqrt{\kappa_1}+\sqrt{\kappa_2}\right)^2-1\right]\left[\log(\epsilon^{-1})+0.73\,(d+1)\right]\right)^{1/2}\right]^2 \\
&= \mathbf{O}\Big((\kappa_1+\kappa_2)\big[d+\log(\epsilon^{-1})\big]\Big).
\end{aligned}$$

With probability at least $1-4\epsilon$*, for any* $\theta\in\mathbb{R}^d$*, any* $\xi\in\mathbb{R}$*,*

$$-\frac{\widehat{\delta}+\gamma_-}{1+\widehat{\delta}} \le \frac{\overline{N}(\theta,\xi)}{N(\theta,\xi)}-1 \le \frac{\widehat{\delta}+\gamma_+}{(1-\widehat{\delta})_+(1-\gamma_+)_+} \le \frac{\widehat{\delta}+\widehat{\gamma}_+}{(1-\widehat{\delta})_+(1-\widehat{\gamma}_+)_+}$$

so that in particular

$$\left|\frac{\overline{N}(\theta,\xi)}{N(\theta,\xi)}-1\right| \le \frac{\widehat{\delta}+\widehat{\gamma}_+}{(1-\widehat{\delta})_+(1-\widehat{\gamma}_+)_+} = \mathbf{O}_{n\to\infty}\left(\sqrt{\frac{(\kappa_1+\kappa 2)\big[d+\log(\epsilon^{-1})\big]}{n}}\right).$$

As a consequence, the empirical risk minimizer

$$\widehat{\theta} \in \arg\min \sum_{i=1}^{n} \big(Y_i - \langle \theta, X_i \rangle\big)^2$$

is such that with probability at least $1 - 4\epsilon$,

$$\mathbb{E}\big[\big(Y - \langle \widehat{\theta}, X\rangle\big)^2\big] - \mathbb{E}\big[\big(Y - \langle \theta_*, X\rangle\big)^2\big] \leq \frac{\delta^2}{(1-\delta)(1-\delta^2)n} \sum_{i=1}^{n} \big(Y_i - \langle \widehat{\theta}, X_i\rangle\big)^2 \leq \frac{\delta^2}{(1-\delta)^2} \mathbb{E}\big[\big(Y - \langle \theta_*, X\rangle\big)^2\big],$$

where

$$\delta^2 = \left(\frac{\widehat{\delta} + \widehat{\gamma}_+}{(1-\widehat{\delta})(1-\widehat{\gamma}_+)}\right)^2 = \mathbf{O}_{n\to\infty}\left(\frac{(\kappa_1+\kappa_2)\big[d + \log(\epsilon^{-1})\big]}{n}\right).$$

We can also apply Proposition 2.4 on page 15 and work under polynomial moment assumptions. Consider the Gram matrix

$$\widetilde{G} = \mathbb{E}\left[\begin{pmatrix} X \\ -Y \end{pmatrix}\big(X^\top, -Y\big)\right].$$

Remark that

$$\left\|\widetilde{G}^{-1/2}\begin{pmatrix} X \\ -Y \end{pmatrix}\right\|^2 = \sup\Big\{\big(\xi Y - \langle \theta, X\rangle\big)^2 : (\theta,\xi) \in \mathbb{R}^{d+1}, \mathbb{E}\big[\big(\xi Y - \langle \theta, X\rangle\big)^2\big] \leq 1\Big\} \leq \frac{(Y - \langle \theta_*, X\rangle)^2}{\mathbb{E}\big[\big(Y - \langle \theta_*, X\rangle\big)^2\big]} + \|G^{-1/2}X\|^2,$$

according to equation (14) on page 27. Therefore, according to the Minkowski inequality in $\mathbb{L}^{p+1}$,

$$\mathbb{E}\left(\left\|\widetilde{G}^{-1/2}\begin{pmatrix} X \\ -Y \end{pmatrix}\right\|^{2p+2}\right) \leq \left(\frac{\mathbb{E}\big[\big(Y - \langle \theta_*, X\rangle\big)^{2p+2}\big]^{1/(p+1)}}{\mathbb{E}\big[\big(Y - \langle \theta_*, X\rangle\big)^2\big]} + \mathbb{E}\big(\|G^{-1/2}X\|^{2p+2}\big)^{1/(p+1)}\right)^{p+1}.$$

In the same way, with a change of notation,

$$\mathbb{E}\left(\left\|\widetilde{G}^{-1/2}\begin{pmatrix}X\\-Y\end{pmatrix}\right\|^{4p+4}\right) \leq \left(\frac{\mathbb{E}\big[\big(Y-\langle\theta_*,X\rangle\big)^{4p+4}\big]^{1/(2p+2)}}{\mathbb{E}\big[\big(Y-\langle\theta_*,X\rangle\big)^2\big]} + \mathbb{E}\big(\|G^{-1/2}X\|^{4p+4}\big)^{1/(2p+2)}\right)^{2p+2}.$$

Proposition 4.7 *Consider some exponent $p \in]1,2]$. Consider the same hypotheses as in Proposition 4.6 on page 28, except that instead of conditions (15) and (16) on page 27 we assume now that*

$$\mathbb{E}\big[\big(Y-\langle\theta_*,X\rangle\big)^{4p+4}\big] < \infty,$$
$$\text{and } \mathbb{E}\big(\|G^{-1/2}X\|^{4p+4}\big) < \infty.$$

Define

$$\widetilde{\gamma}_+ = \frac{1}{p+1}\left(\frac{2\big[\log(\epsilon^{-1})+0.73\,(d+1)\big]}{3\big[(\sqrt{\kappa_1}+\sqrt{\kappa_2})^2-1\big)n}\right)^{p/2}(1+\widehat{\delta})^{p+1}$$
$$\times\left[\left(\frac{\mathbb{E}\big[\big(Y-\langle\theta_*,X\rangle\big)^{2p+2}\big]^{1/(p+1)}}{\mathbb{E}\big[\big(Y-\langle\theta_*,X\rangle\big)^2\big]}+\mathbb{E}\big(\|G^{-1/2}X\|^{2p+2}\big)^{1/(p+1)}\right)^{p+1}\right.$$
$$\left.+\frac{1}{\sqrt{n\epsilon}}\left(\frac{\mathbb{E}\big[\big(Y-\langle\theta_*,X\rangle\big)^{4p+4}\big]^{1/(2p+2)}}{\mathbb{E}\big[\big(Y-\langle\theta_*,X\rangle\big)^2\big]}+\mathbb{E}\big(\|G^{-1/2}X\|^{4p+4}\big)^{1/(2p+2)}\right)^{p+1}\right],$$

where κ_1 and κ_2 are as in Lemma 2.5 on page 16. Under the same condition on n and ϵ as in Proposition 4.6, with probability at least $1-3\epsilon$, for any $(\theta,\xi)\in\mathbb{R}^{d+1}$,

$$\left|\frac{\overline{N}(\theta,\xi)}{N(\theta,\xi)}-1\right| \leq \frac{\widehat{\delta}+\widetilde{\gamma}_+}{(1-\widehat{\delta})_+(1-\widetilde{\gamma}_+)_+},$$

so that in particular the empirical risk minimizer is such that with probability at least $1-3\epsilon$,

$$\mathbb{E}\big[\big(Y-\langle\widehat{\theta},X\rangle\big)^2\big]-\mathbb{E}\big[\big(Y-\langle\theta_*,X\rangle\big)^2\big] \leq \frac{\delta^2}{(1-\delta)(1-\delta^2)n}\sum_{i=1}^n\big(Y_i-\langle\widehat{\theta},X_i\rangle\big)^2$$
$$\leq \frac{\delta^2}{(1-\delta)^2}\mathbb{E}\big[\big(Y-\langle\theta_*,X\rangle\big)^2\big],$$

where

$$\delta^2 = \left(\frac{\widehat{\delta} + \widetilde{\gamma}_+}{(1-\widehat{\delta})(1-\widetilde{\gamma}_+)}\right)^2 = \mathbf{O}_{n\to\infty}\big(\widehat{\delta}^2\big) = \mathbf{O}_{n\to\infty}\left(\frac{(\kappa_1+\kappa_2)\big[d+\log(\epsilon^{-1})\big]}{n}\right).$$

We can also apply Proposition 2.6 on page 17 to get weaker moment assumptions at the price of a worse non-asymptotic bound, but still with the same leading term when the sample size n goes to infinity.

Proposition 4.8 *Consider some exponents $p \in]1,2]$ and $q \in]1,2[$. Make the same hypotheses as in Proposition 4.6 on page 28, except that instead of conditions* (15) *and* (16) *on page 27, we assume that*

$$\mathbb{E}\big[\big(Y - \langle\theta_*, X\rangle\big)^{2q(p+1)}\big] < \infty \text{ and } \mathbb{E}\big(\|G^{-1/2}X\|^{2q(p+1)}\big) < \infty.$$

Define

$$\begin{aligned}\widetilde{\gamma}_+ = \frac{1}{p+1}\left(\frac{2\big[\log(\epsilon^{-1}) + 0.73(d+1)\big]}{\big[(\sqrt{\kappa_1}+\sqrt{\kappa_2})^2 - 1\big)n}\right)^{p/2}\big(1+\widehat{\delta}\big)^{p+1}\\ \times\left[\left(\frac{\mathbb{E}\big[\big(Y-\langle\theta_*,X\rangle\big)^{2p+2}\big]^{1/(p+1)}}{\mathbb{E}\big[\big(Y-\langle\theta_*,X\rangle\big)^2\big]} + \mathbb{E}\big(\|G^{-1/2}X\|^{2p+2}\big)^{1/(p+1)}\right)^{p+1}\right.\\ + \frac{C_q}{\epsilon^{1/q}n^{1-1/q}}\left(\frac{\mathbb{E}\big[\big(Y-\langle\theta_*,X\rangle\big)^{2q(p+1)}\big]^{q^{-1}(p+1)^{-1}}}{\mathbb{E}\big[\big(Y-\langle\theta_*,X\rangle\big)^2\big]}\right.\\ \left.\left.+\,\mathbb{E}\Big(\|G^{-1/2}X\|^{2q(p+1)}\Big)^{q^{-1}(p+1)^{-1}}\right)^{p+1}\right],\end{aligned}$$

where the constant C_q is defined in Lemma 2.5 on page 16 and κ_1 and κ_2 are defined in Lemma 4.3 on page 25. Under the same conditions on n and ϵ as in Proposition 4.6 on page 28, with probability at least $1-4\epsilon$, for any $(\theta,\xi)\in\mathbb{R}^{d+1}$,

$$\left|\frac{\overline{N}(\theta,\xi)}{N(\theta,\xi)} - 1\right| \le \frac{\widehat{\delta}+\widetilde{\gamma}_+}{(1-\widehat{\delta})_+(1-\widetilde{\gamma}_+)_+},$$

so that the empirical risk minimizer is such that with probability at least $1-4\epsilon$,

$$\mathbb{E}\big[\big(Y-\langle\widehat{\theta},X\rangle\big)^2\big] - \mathbb{E}\big[\big(Y-\langle\theta_*,X\rangle\big)^2\big]$$

$$\leq \frac{\delta^2}{(1-\delta)_+(1-\delta^2)_+ n} \sum_{i=1}^n \big(Y_i - \langle \widehat{\theta}, X_i \rangle\big)^2$$
$$\leq \frac{\delta^2}{(1-\delta)^2} \mathbb{E}\big[\big(Y - \langle \theta_*, X \rangle\big)^2\big],$$

where

$$\delta^2 = \left(\frac{\widehat{\delta} + \widetilde{\gamma}_+}{(1-\widehat{\delta})_+(1-\widetilde{\gamma}_+)_+}\right)^2 = \mathbf{O}_{n\to\infty}\big(\widehat{\delta}^2\big)$$
$$= \mathbf{O}_{n\to\infty}\left(\frac{(\kappa_1+\kappa_2)\big[d + \log(\epsilon^{-1})\big]}{n}\right).$$

Remark that in these last two propositions, the hypotheses on the design X and the noise $Y - \langle \theta_*, X \rangle$ are weaker than in Proposition 4.6, since they involve only polynomial moment assumptions. As a counterpart, the dependence in the confidence parameter ϵ is worse, since we have a factor $1/\sqrt{n\epsilon}$ in Proposition 4.7 and $\epsilon^{-1/q} n^{-(1-1/q)}$ in Proposition 4.8 that precludes the use of a confidence level $1 - 3\epsilon$ (resp. $1 - 4\epsilon$) much higher than $1 - 1/n$ (resp. $1 - n^{-(q-1)}$). Nevertheless, this factor $1/\sqrt{n\epsilon}$ (resp. $\epsilon^{-1/q} n^{-(1-1/q)}$) appears only in second order terms, so that we still get an asymptotic upper bound (expressed here in big **O** notation) where the dependence on the confidence parameter is proportional to $\log(\epsilon^{-1})$.

The results in this section are stronger than those obtained in [5]. In particular, for the robust estimator, we are not limited to considering the minimization in θ in a bounded subset Θ of $\mathbb{R}^d$, as in Theorem 3.1 of [5], where some of the constants involved in the bound depend on the fact that Θ is bounded. Also, we provide non-asymptotic bounds for the empirical risk minimizer, whereas in [5] the corresponding result, Theorem 2.1, is asymptotic, since it is satisfied only for large enough sample sizes n, without an explicit condition on n. Moreover the constant B appearing in Theorem 2.2 of [5] is in fact necessarily dependent on the dimension and more precisely not smaller than d, as mentioned to us by Guillaume Lécué, whom we are grateful for pointing out this mistake (the comments we made after Theorem 2.2 of [5] about the size of the constant B are in fact false, we apologize for this error).

It is also interesting to compare our results with those of [15]. In particular, it is relevant to compare Proposition 4.8 above with Theorem 1.3 of [15]. Our hypotheses are only slightly stronger than those of [15], since we require that $\|G^{-1/2}X\|$ belongs to a little more than $\mathbb{L}^4$, namely to $\mathbb{L}^{2q(p+1)}$, where $q > 1$ and $p > 1$ can be taken arbitrarily close to 1, so that $2q(p+1)$

can also be made as close to 4 as desired, if one is willing to accept a larger second order term $\widetilde{\gamma}_+$. Indeed, the coefficient θ_0 appearing in [15] is related to the kurtosis coefficient κ_1 by the relation $\theta_0 = \kappa_1^{1/4}$. Moreover, as proved in Lemma A.7 on page 51

$$\mathbb{E}\big(\|G^{-1/2}X\|^4\big) \leq d\kappa_1,$$

so that when κ_1 (or θ_0) is finite, $\|G^{-1/2}X\|$ belongs to $\mathbb{L}^4$. So we ask for a little higher moment, and we get as a reward a better bound, since our bound still has a subexponential first order term (meaning that the dependence in the confidence parameter ϵ is in $\log(\epsilon^{-1})$), whereas the bound obtained in [15] is

$$\mathbf{O}_{n\to\infty}\left[\frac{d\kappa_1^3\kappa_2^{1/2}}{n\epsilon}\mathbb{E}\big[\big(Y-\langle\theta_*,X\rangle\big)^2\big]\right],$$

when cast into our notation. For instance, if we choose $p = q = 3/2$ in our bound, we have to assume that $\mathbb{E}\big(\|G^{-1/2}X\|^{15/2}\big) < \infty$, but we can take $\epsilon = n^{-1/2}$ and still have a second order term in $n^{-3/2}$, whereas the first order term is in n^{-1}. As a comparison, if we choose $\epsilon = n^{-1/2}$, the bound from [15] is no more in n^{-1}, but in $n^{-1/2}$.

We would like also to remark that in some situations, κ_1, that is equal to θ_0^4 in [15], may in fact depend on d. This is for instance the case in uniform histogram regression, where

$$X = \big(X_k, k = 1, \dots d\big) = \big[g_k(U), k = 1, \dots, d\big] \tag{17}$$

the random variable U being distributed according to the uniform distribution in the unit interval $[0, 1]$ and the functions g_k being defined as

$$g_k(u) = \mathbb{1}\big[\,(k-1)/d \leq u < k/d\,\big]. \tag{18}$$

In this case $\mathbb{E}\big[g_k(U)^p\big] = 1/d$ for all values of the exponent p, so that necessarily $\kappa_1 \geq \mathbb{E}\big[g_k(U)^4\big]/\mathbb{E}\big[g_k(U)^2\big]^2 = d$ (in fact, it is easy to see that $\kappa_1 = d$ in this case). The same unfavourable scaling appears in all local bases of the wavelet type. Indeed, in the context of functional regression of Y by $\sum_{k=1}^d \theta_k g_k(U)$, where U is uniform in the unit interval, rescaling a regression function g by a scale factor λ will impact its $\mathbb{L}^4/\mathbb{L}^2$ ratio according to the formula

$$\frac{\mathbb{E}\big[g(\lambda U)^4\big]}{\mathbb{E}\big[g(\lambda U)^2\big]^2} = \lambda\frac{\mathbb{E}\big[g(U)^4\big]}{\mathbb{E}\big[g(U)^2\big]^2}.$$

4.4. Some lower bound for the empirical risk minimizer. We will show in this section that the order of magnitude of the previous upper bound cannot be improved in the worst case, except for the values of the numerical constants.

Proposition 4.9 *Consider a sample $(\widetilde{X}_1,\widetilde{Y}_1),\dots,(\widetilde{X}_n,\widetilde{Y}_n)$ made of n independent copies of the couple of random variables $(\widetilde{X},\widetilde{Y})\in\mathbb{R}^d\times\mathbb{R}$. Assume that for some $\delta\in]0,1[$, some $\epsilon\in]0,1/2[$ and some $n_\epsilon\in\mathbb{N}$, for any $n\geq n_\epsilon$, with probability at least $1-2\epsilon$, for any $\theta\in\mathbb{R}^d$,*

$$\left|\frac{\frac{1}{n}\sum_{i=1}^n\langle\theta,\widetilde{X}_i\rangle^2}{\mathbb{E}\big(\langle\theta,\widetilde{X}\rangle^2\big)}-1\right|\leq\delta, \tag{19}$$

where by convention $0/0=1$ and $z/0=\infty$ for $z>0$. Assume that

$$\widetilde{Y}=\langle\theta_*,\widetilde{X}\rangle+\eta,$$

where $\mathbb{P}_{\eta|\widetilde{X}}=\mathcal{N}(0,\sigma^2)$, with $\sigma>0$, so that the noise η is a Gaussian noise independent from $\widetilde{X}$. Assume also that

$$G=\mathbb{E}\big(\widetilde{X}\widetilde{X}^\top\big)$$

is of full rank d. Consider any empirical risk minimizer

$$\widetilde{\theta}\in\arg\min_{\theta\in\mathbb{R}^d}\sum_{i=1}^n\big(\widetilde{Y}_i-\langle\theta,\widetilde{X}_i\rangle\big)^2.$$

(There may be more than one when $\frac{1}{n}\sum_{i=1}^n\widetilde{X}_i\widetilde{X}_i^\top$ is not of full rank.)

For any $n\geq n_\epsilon$, with probability at least $1-3\epsilon$,

$$\mathbb{E}\big[\big(\widetilde{Y}-\langle\widetilde{\theta},\widetilde{X}\rangle\big)^2\big]-\mathbb{E}\big[\big(\widetilde{Y}-\langle\theta_*,\widetilde{X}\rangle\big)^2\big]\geq\frac{\big[d\log(2)-2\log(\epsilon^{-1})\big]\sigma^2}{(1+\delta)\,n}.$$

Consider now a Bernoulli random variable ξ of parameter $p\in]0,1]$, independent from $(\widetilde{X},\widetilde{Y})$, so that

$$\mathbb{P}\big(\xi=1\,|\,(\widetilde{X},\widetilde{Y})\big)=1-\mathbb{P}\big(\xi=0\,|\,(\widetilde{X},\widetilde{Y})\big)=p.$$

Define the censored couple of random variables

$$\big(X,Y\big)=\big(\xi\widetilde{X},\xi\widetilde{Y}\big).$$

Consider a sample $(X_1, Y_1), \dots, (X_n, Y_n)$ made of n independent copies of (X, Y). Consider any empirical risk minimizer of the censored data

$$\widehat{\theta} \in \arg\min_{\theta \in \mathbb{R}^d} \sum_{i=1}^{n} \big(Y_i - \langle \theta, X_i \rangle\big)^2.$$

Define κ_1 and κ_2 as in Lemma 4.3 on page 25. Define in the same way

$$\widetilde{\kappa}_1 = \sup\Big\{ \mathbb{E}\big(\langle \theta, \widetilde{X}\rangle^4\big) \,:\, \theta \in \mathbb{R}^d,\ \mathbb{E}\big(\langle \theta, \widetilde{X}\rangle^2\big) \le 1 \Big\},$$

$$\widetilde{\kappa}_2 = \frac{\mathbb{E}\big[\big(\widetilde{Y} - \langle \theta_*, \widetilde{X}\rangle\big)^4\big]}{\mathbb{E}\big[\big(\widetilde{Y} - \langle \theta_*, \widetilde{X}\rangle\big)^2\big]^2} = 3 \text{ (since } \eta \text{ is Gaussian).}$$

Remark that $\kappa_1 = \widetilde{\kappa}_1/p$ and $\kappa_2 = \widetilde{\kappa}_2/p$, so that in particular $p^{-1} = (\kappa_1 + \kappa_2)/(\widetilde{\kappa}_1 + \widetilde{\kappa}_2)$. For any

$$n \ge \frac{4(\kappa_1 + \kappa_2) \max\big\{n_\epsilon, 2\log(\epsilon^{-1})\big\}}{\widetilde{\kappa}_1 + \widetilde{\kappa}_2}$$

with probability at least $1 - 5\epsilon$,

$$\mathbb{E}\big[\big(Y - \langle \widehat{\theta}, X\rangle\big)^2\big] - \mathbb{E}\big[\big(Y - \langle \theta_*, X\rangle\big)^2\big] \ge \frac{4\big[d\log(2) - 2\log(\epsilon^{-1})\big](\kappa_1 + \kappa_2)}{7(1+\delta)(\widetilde{\kappa}_1 + \widetilde{\kappa}_2)\, n} \mathbb{E}\big[\big(Y - \langle \theta_*, X\rangle\big)^2\big]. \quad (20)$$

Therefore, the first order term dependence in $\kappa_1 + \kappa_2$ in Propositions 4.7 on page 30 and 4.8 cannot be improved without further assumptions, since by varying the value of p^{-1}, we can make $\kappa_1 + \kappa_2$ arbitrarily large in this lower bound.

In the case when $\widetilde{X}$ is a Gaussian vector, $\widetilde{\kappa}_1 + \widetilde{\kappa}_2 = 6$ and we can take for example

$$n_\epsilon = 8.4\big[0.73\, d + \log(\epsilon^{-1})\big]\big[1.4\, d + 4\log(\epsilon^{-1}) + 26.3\big]^2.$$

When $n \ge n_\epsilon$, we can deduce from Proposition 2.3 on page 14 that equation (19) *on page 34 holds with $\delta = 8/9$.*

So in the case of a censored Gaussian design, for any $n \ge \dfrac{2}{3}(\kappa_1 + \kappa_2) n_\epsilon$, with probability at least $1 - 5\epsilon$,

$$\mathbb{E}\big[\big(Y - \langle \widehat{\theta}, X\rangle\big)^2\big] - \mathbb{E}\big[\big(Y - \langle \theta_*.X\rangle\big)^2\big] \ge \frac{35}{1000} \times \frac{\big[d - 3\log(\epsilon^{-1})\big](\kappa_1 + \kappa_2)}{n} \mathbb{E}\big[\big(Y - \langle \theta_*, X\rangle\big)^2\big].$$

Note that the result given in the end of this proposition is not the best possible and could be improved by using PAC-Bayes bounds specific to the Gaussian case, resulting in a smaller n_ϵ with a better dependence in the dimension d. However, since we will give another lower bound further on, we thought it wiser not to delve into these technicalities. The proof of this proposition is given in the appendix.

4.5. Exact convergence rate for the empirical risk minimizer. The results of the previous section show that the convergence speed of the empirical risk mimimizer cannot be $\mathbf{O}_{n\to n}\Big(R(\theta_*)d/n\Big)$ in the worst case. In this section, we clarify the situation by showing under mild conditions that

$$\mathbb{E}\Big[\min\big\{R(\widehat{\theta}) - R(\theta_*), C\exp\big(n^{2-q}\big)\big\}\Big] \underset{n\to\infty}{\sim} C/n,$$

where we recall that $R(\theta) = \mathbb{E}\big[\big(Y - \langle\theta, X\rangle\big)^2\big]$ and where the exact constant C is given by

$$C = \mathbb{E}\Big[\big(Y - \langle\theta_*, X\rangle\big)^2\big\|G^{-1/2}X\big\|^2\Big].$$

The form of this constant shows that we get a $R(\theta_*)\,d/n$ convergence rate when the noise $Y - \langle\theta_*, X\rangle$ is independent from X (and G is of full rank), and that it can be larger or smaller otherwise.

Proposition 4.10 *Consider a sample $(X_1, Y_1), \dots, (X_n, Y_n)$ made of n independent copies of the couple of random variables $(X, Y) \in \mathbb{R}^d \times \mathbb{R}$. Let $G = \mathbb{E}\big(XX^\top\big)$ be the Gram matrix of the design X and G^{-1} its pseudo-inverse. Assume that*

$$\kappa = \sup\Big\{\mathbb{E}\big(\langle\theta, X\rangle^4\big) \,:\, \theta\in\mathbb{R}^d, \mathbb{E}\big(\langle\theta, X\rangle^2\big) \le 1\Big\} < \infty$$

and that

$$C = \mathbb{E}\Big[\big(Y - \langle\theta_*, X\rangle\big)^2\big\|G^{-1/2}X\big\|^2\Big] < \infty,$$

where

$$\theta_* \in \arg\min_{\theta\in\mathbb{R}^d} \mathbb{E}\big[\big(Y - \langle\theta, X\rangle\big)^2\big]$$

is some optimal regression parameter. Define μ as in Proposition 1.1 on page 7 and γ_- and $\widehat{\delta}$ as in Proposition 2.2 on page 11. Consider any empirical risk minimizer

$$\widehat{\theta} \in \arg\min_{\theta\in\mathbb{R}^d} \sum_{i=1}^n \big(Y_i - \langle\theta, X_i\rangle\big)^2.$$

For any $n \ge n_\epsilon = \mathbf{O}\big(\kappa\big[d + \log(\epsilon^{-1})\big]\big)$ given by equation (3) on page 7, there is an event Ω of probability at least $1-\epsilon$, such that the excess risk conditional to Ω satisfies

$$\frac{n}{C}\,\mathbb{E}\Big[R(\widehat{\theta}\,) - R(\theta_*)\,\Big|\,\Omega\Big] \le \frac{(1+\widehat{\delta})^2}{(1-\gamma_-)^2\,\mathbb{P}(\Omega)} \le \left(1 + \mathbf{O}\left(\sqrt{\frac{\kappa\big[d+\log\big(\epsilon^{-1}\big)\big]}{n}}\,\right)\right)(1-\epsilon)^{-1}.$$

Consequently, for any ϵ and n satisfying equation (3) on page 7, for any positive constant $M \in \mathbb{R}_+$,

$$\mathbb{E}\Big[\min\big\{R(\widehat{\theta}\,) - R(\theta_*), M\big\}\Big] \le \frac{(1+\widehat{\delta})^2\,C}{(1-\gamma_-)^2\,n} + \epsilon M.$$

Taking $\epsilon = \dfrac{C}{Mn^2}$, we see that equation (3) on page 7 is satisfied for $n > \mathbf{O}\Big(\kappa\Big[d + \log\big[\kappa M/C\big]\Big]\Big)$, and that

$$\frac{n}{C}\,\mathbb{E}\Big[\min\big\{R(\widehat{\theta}) - R(\theta_*), M\big\}\Big] \le \left(1 + \mathbf{O}\left(\sqrt{\frac{\kappa\big[d+\log\big(Mn/C\big)\big]}{n}}\,\right)\right).$$

Assume now moreover that $Y = \langle \theta_*, X\rangle + \eta$, where $\theta_* \in \mathbb{R}^d$ and η is a random variable independent from X such that $\mathbb{E}(\eta) = 0$ and $0 < \mathbb{E}(\eta^2) = \sigma^2 < \infty$. In this case, for $n \ge \mathbf{O}\Big(\kappa\big[d + \log\big(\kappa M/\sigma^2\big)\big]\Big)$,

$$\mathbb{E}\Big[\min\big\{R(\widehat{\theta}\,) - R(\theta_*), M\big\}\Big] \le \left(1 + \mathbf{O}\left(\sqrt{\frac{\kappa\big[d+\log\big(Mn/\sigma^2\big)\big]}{n}}\,\right)\right)\frac{\sigma^2 d}{n}.$$

Moreover, when $n \ge \mathbf{O}\Big(\kappa\big[d + \log(\epsilon^{-1})\big]\Big)$ satisfies equation (3) on page 7, with probability at least $1-\epsilon$,

$$\mathbb{E}\big[R(\widehat{\theta}\,) - R(\theta_*)\,\big|\,X_1,\dots,X_n\big] \le \frac{(1+\widehat{\delta})\,d\,\sigma^2}{(1-\gamma_-)\,n} = \left(1 + \mathbf{O}\left(\sqrt{\frac{\kappa\big[d+\log(\epsilon^{-1})\big]}{n}}\,\right)\right)\frac{d\,\sigma^2}{n}.$$

Remark that the last statement of this proposition is about the random design risk

$$R(\theta) = \mathbb{E}\big[\big(Y - \langle\theta, X\rangle\big)^2\big]$$

and not about the weaker fixed design risk

$$\mathbb{E}\left(\frac{1}{n}\sum_{i=1}^{n}\big(Y_i - \langle\theta, X_i\rangle\big)^2 \,\Big|\, X_1, \dots, X_n\right).$$

The proof of Proposition 4.10 is presented in the appendix.

The next proposition states an upper bound with large probability when the noise is Gaussian and independent from the design X.

Proposition 4.11 *Make the same assumptions as in the end of the previous proposition. Assume moreover that the noise η is Gaussian (but not necessarily the design X). For any $n > n_\epsilon = \mathbf{O}\big(\kappa\big[d + \log(\epsilon^{-1})\big]\big)$ given by equation* (3) *on page 7, with probability at least $1 - 2\epsilon$,*

$$\begin{aligned} R\big(\widehat{\theta}\big) - R(\theta_*) &\le \frac{(1+\widehat{\delta})\,\sigma^2}{(1-\gamma_-)\,n}\big[2\,d\,\log(2) + 4\log(\epsilon^{-1})\big] \\ &= \left(1 + \mathbf{O}\left(\sqrt{\frac{\kappa\big[d+\log\big(\epsilon^{-1}\big)\big]}{n}}\,\right)\right)\frac{\sigma^2\big[2\,d\log(2) + 4\log(\epsilon^{-1})\big]}{n}, \end{aligned}$$

where it is interesting to remind that $2\log(2) \le 1.4$.

Proof. Using the same notation as in the proof of Proposition 4.10 on page 36,

$$\mathbb{P}_{\overline{G}^{-1}W|X_1,\dots,X_n} = \mathcal{N}\left(0, \frac{\sigma^2}{n}\overline{G}^{-1}\overline{G}\right),$$

where $\overline{G}^{-1}\overline{G}$ is the orthogonal projection on $\mathbf{Im}(\overline{G})$. Therefore, we can modify the end of the proof of Proposition 4.10, using a Chernoff deviation bound for Gaussian vectors. □

Now, let us close this section with the corresponding lower bound, to get the announced exact convergence rate in expectation of the excess risk of the empirical risk minimizer.

Proposition 4.12 *Consider a sample $(X_1, Y_1), \dots, (X_n, Y_n)$ made of n independent copies of the couple of random variables $(X, Y) \in \mathbb{R}^d \times \mathbb{R}$. Choose two exponents $p, q \in [1, 2]$. Assume that*

$$\begin{aligned} &\kappa = \sup\big\{\mathbb{E}\big(\langle\theta, X\rangle^4\big) \,:\, \theta \in \mathbb{R}^d, \mathbb{E}\big(\langle\theta, X\rangle^2\big) \le 1\big\} < \infty, \\ &C = \mathbb{E}\Big[\big(Y - \langle\theta_*, X\rangle\big)^2\big\|G^{-1/2}X\big\|^2\Big] < \infty, \\ &C > 0, \end{aligned}$$

$$\kappa' = \frac{\mathbb{E}\Big[\big(Y - \langle \theta_*, X\rangle\big)^4 \big\|G^{-1/2}X\big\|^4\Big]}{\mathbb{E}\Big[\big(Y - \langle \theta_*, X\rangle\big)^2 \big\|G^{-1/2}X\big\|^2\Big]^2} < \infty,$$
$$\mathbb{E}\Big(\big\|G^{-1/2}X\big\|^{2q(p+1)}\Big) < \infty.$$

Under these hypotheses, we will give a technical meaning, in two different ways, to the fact that C/n is the exact convergence rate of the excess risk $R(\widehat{\theta}) - R(\theta_*)$ of the empiricial risk minimizer

$$\widehat{\theta} \in \arg\min_{\theta\in\mathbb{R}^d} \sum_{i=1}^n \big(Y_i - \langle \theta, X_i\rangle\big)^2.$$

Define μ as in Proposition 1.1 on page 7, $\widehat{\delta} = \mu/(1-2\mu)$ and define

$$\widetilde{\gamma}_+ = \mathbf{O}\left(\left(\frac{\log(\epsilon^{-1}) + d}{\kappa n}\right)^{p/2}\left[\mathbb{E}\big(\big\|G^{-1/2}X\big\|^{2(p+1)} + \frac{\mathbb{E}\big(\big\|G^{-1/2}X\big\|^{2q(p+1)}\big)^{1/q}}{\epsilon^{1/q} n^{1-1/q}}\right]\right)$$

as in Proposition 2.6 on page 17.

For any $n \geq n_\epsilon = \mathbf{O}\big(\kappa\big[d + \log(\epsilon^{-1})\big]\big)$ given by equation (3) on page 7, there is an event Ω of probability at least $1 - 2\epsilon$ such that

$$\begin{aligned}(1-\widetilde{\gamma}_+)_+^2(1-\widehat{\delta})_+^2\left[1 - \sqrt{6}\left(1 + \frac{\kappa'-3}{3n}\right)^{1/2}\epsilon^{1/2}\right] \times \frac{C}{n}& \\ \leq \mathbb{E}\Big(\big[R(\widehat{\theta}) - R(\theta_*)\big]\mathbb{1}_\Omega\Big) \leq \mathbb{E}\Big(R(\widehat{\theta} - R(\theta_*)\,\big|\,\Omega\Big)& \\ \leq \frac{(1+\widehat{\delta})^2}{(1-\gamma_-)^2(1-2\epsilon)} \times \frac{C}{n}.& \quad (21)\end{aligned}$$

Remark that when $p > 1$, this gives

$$\left|\frac{n}{C}\,\mathbb{E}\Big(R(\widehat{\theta}) - R(\theta_*)\,\big|\,\Omega\Big) - 1\right| \leq \mathbf{O}_{n\to\infty}\left(\sqrt{\frac{\kappa\big[d + \log(\epsilon^{-1})\big]}{n}} + \epsilon^{1/2}\right),$$
$$\textit{and}\quad \left|\frac{n}{C}\,\mathbb{E}\Big(\big[R(\widehat{\theta}) - R(\theta_*)\big]\mathbb{1}_\Omega\Big) - 1\right| \leq \mathbf{O}_{n\to\infty}\left(\sqrt{\frac{\kappa\big[d + \log(\epsilon^{-1})\big]}{n}} + \epsilon^{1/2}\right).$$

For any $M > 0$, any ϵ satisfying equation (3) on page 7,

$$\frac{n}{C}\,\mathbb{E}\Big[\min\big\{R(\widehat{\theta}) - R(\theta_*), M\big\}\Big]$$

$$\geq (1-\widehat{\delta})_+^2(1-\widetilde{\gamma}_+)_+^2\Bigg[1-\frac{3C}{4Mn}\bigg(1+\frac{\kappa'-3}{3n}\bigg) - \sqrt{6}\bigg(1+\frac{\kappa'-3}{3n}\bigg)^{1/2}\epsilon^{1/2}\Bigg]. \quad (22)$$

If we assume that $q>1$ and bind ϵ and n by the relation $\epsilon = n^{-(q-1)}$, equation (3) on page 7 is satisfied when $n \geq \mathbf{O}\big(\kappa\big[d+\log(\kappa)\big]\big)$ and

$$\widetilde{\gamma}_+ = \mathbf{O}\Bigg(\bigg(\frac{\log(n)+d}{\kappa n}\bigg)^{p/2}\mathbb{E}\Big(\big\|G^{-1/2}X\big\|^{2q(p+1)}\Big)^{1/q}\Bigg),$$

so that

$$\begin{aligned}\frac{n}{C}\mathbb{E}\Big[\min\big\{R(\widehat{\theta})-R(\theta_*),M\big\}\Big] &\geq 1-\mathbf{O}\Bigg(\sqrt{\frac{\kappa\big[d+\log(n)\big]}{n}}+\bigg(\frac{\log(n)+d}{\kappa n}\bigg)^{p/2}\mathbb{E}\Big(\big\|G^{-1/2}X\big\|^{2q(p+1)}\Big)^{1/q} \\ &\qquad + \frac{(1+\kappa'/n)C}{Mn}+\big(1+\kappa'/n\big)^{1/2}n^{-(q-1)/2}\Bigg) \\ &\underset{when\ q<2}{=} 1-\mathbf{O}_{n\to\infty}\Big(n^{-(q-1)/2}\Big).\end{aligned}$$

Combining this result with the reverse bound of Proposition 4.10 on page 36 gives, in the case when $q\in]1,2[$,

$$\bigg|\frac{n}{C}\,\mathbb{E}\Big[\min\big\{R(\widehat{\theta})-R(\theta_*),M\big\}\Big]-1\bigg| \leq \mathbf{O}_{n\to\infty}\big(n^{-(q-1)/2}\big).$$

Looking at the non-asymptotic bounds, we see that we can bind M to n by the relation $M=C\exp\big(n^{2-q}\big)$ and still get when $q\in]1,2[$,

$$\bigg|\frac{n}{C}\mathbb{E}\Big[\min\big\{R(\widehat{\theta})-R(\theta_*),\,C\exp\big(n^{2-q}\big)\big\}\Big]-1\bigg| \leq \mathbf{O}_{n\to\infty}\big(n^{-(q-1)/2}\big). \quad (23)$$

This means that we can threshold $R(\widehat{\theta})-R(\theta_*)$ at a very high level, in the sens that the threshold is reached with a very small probability when the sample size is large, since from Markov's inequality,

$$\begin{aligned}\mathbb{P}\Big[R(\widehat{\theta})-R(\theta_*)\geq C\exp(n^{2-q})\Big] &\leq \frac{\mathbb{E}\Big[\min\big\{R(\widehat{\theta})-R(\theta_*),C\exp\big(n^{2-q}\big)\big\}\Big]}{C\exp\big(n^{2-q}\big)} \\ &\leq \mathbf{O}_{n\to\infty}\bigg(\frac{1}{n\exp\big(n^{2-q}\big)}\bigg).\end{aligned}$$

The proof of this proposition is given in appendix.

The exact rate C/n can be used to construct another lower bound, this time for the expected excess risk, to complement the lower bound on the deviations of the excess risk given in Proposition 4.9 on page 34. For this, we want to describe a case where C is much larger than $d\,R(\theta_*)/n$. Consider $\eta \in \{-1,+1\}$, a Rademacher random variable independent from X. This means more precisely that

$$\mathbb{P}(\eta = +1 \,|\, X) = \mathbb{P}(\eta = -1 \,|\, X) = 1/2.$$

Define for some $\theta_* \in \mathbb{R}^d$

$$Y = \langle \theta_*, X\rangle + \eta \|G^{-1/2}X\|.$$

Let κ_1 and κ_2 be defined as in Lemma 4.3 on page 25. In this case,

$$\begin{aligned} C = \mathbb{E}\Big(\big(Y - \langle \theta_*, X\rangle\big)^4\Big) &= \kappa_2 \mathbb{E}\Big(\big(Y - \langle \theta_*, X\rangle\big)^2\Big)^2 \\ &= \kappa_2 \mathbb{E}\Big(\big(Y - \langle \theta_*, X\rangle\big)^2\Big)\mathbb{E}\Big(\big\|G^{-1/2}X\big\|^2\Big) = \kappa_2 R(\theta_*) d. \end{aligned}$$

We can add some hypotheses on the structure of X to make sure that κ_2 and $\kappa_1 + \kappa_2$ are of the same order of magnitude. Consider a centered Gaussian vector $W \in \mathbb{R}^d$ such that $W \sim \mathcal{N}(0, G)$ and an independent non-negative real valued random variable ρ such that $\mathbb{E}(\rho^2) = 1$. This implies that $G = \mathbb{E}(XX^\top)$ is the Gram matrix of X as well as of W. Since

$$\begin{aligned} \mathbb{E}\big(\langle \theta, X\rangle^4\big) &= \mathbb{E}\big(\rho^4\big)\mathbb{E}\big(\langle \theta, W\rangle^4\big) \\ &= 3\,\mathbb{E}\big(\rho^4\big)\mathbb{E}\big(\langle \theta, W\rangle^2\big)^2 = 3\,\mathbb{E}\big(\rho^4\big)\mathbb{E}\big(\langle \theta, X\rangle^2\big)^2, \end{aligned}$$

$\kappa_1 = 3\mathbb{E}\big(\rho^4\big)$. On the other hand

$$\begin{aligned} \mathbb{E}\Big(\big(Y - \langle \theta_*, X\rangle\big)^4\Big) &= \mathbb{E}\Big(\big\|G^{-1/2}X\big\|^4\Big) = \mathbb{E}\big(\rho^4\big)\mathbb{E}\Big(\big\|G^{-1/2}W\big\|^4\Big) \\ &= \mathbb{E}\big(\rho^4\big)\big[3d + d(d-1)\big] = \Big(1 + \frac{2}{d}\Big)\mathbb{E}\big(\rho^4\big)\mathbb{E}\Big(\big\|G^{-1/2}X\big\|^2\Big)^2 \\ &= \Big(1 + \frac{2}{d}\Big)\mathbb{E}\big(\rho^4\big)\mathbb{E}\Big(\big(Y - \langle \theta_*, X\rangle\big)^2\Big)^2, \end{aligned}$$

so that $\kappa_2 = \Big(1 + \dfrac{2}{d}\Big)\mathbb{E}\big(\rho^4\big)$ and

$$C = \frac{1 + 2/d}{4 + 2/d}\big(\kappa_1 + \kappa_2\big) R(\theta_*) d.$$

If we assume that $\mathbb{E}\big(\rho^8\big) < \infty$, then the hypotheses of Proposition 4.12 are fulfilled for $p = 1$ and any $q \in [1, 2[$, so that for any $q < 2$,

$$\mathbb{E}\Big(R(\widehat{\theta}) - R(\theta_*)\Big) \geq \mathbb{E}\Big(\min\big\{R(\widehat{\theta}) - R(\theta_*), C\exp\big(n^{2-q}\big)\big\}\Big)$$
$$\geq \Big(1 - \mathbf{O}\big(n^{-(q-1)/2}\big)\Big)\frac{(1+2/d)\big(\kappa_1+\kappa_2\big)dR(\theta_*)}{(4+2/d)n}.$$

This provides a lower bound for the expected excess risk that complements Proposition 4.9 on page 34.

Here, we just described a case where C can be arbitrarily larger than $dR(\theta_*)$, since $\kappa_1 + \kappa_2 = \big(4 + 2/d\big)\mathbb{E}\big(\rho^4\big)$ can be arbitrarily large. It is also interesting to remark that C can be arbitrarily smaller than $d\,R(\theta_*)$. Let us describe such a situation. Let U be a uniform random vector on the unit sphere of $\mathbb{R}^d$, so that $\mathbb{E}\big(UU^\top\big) = d^{-1}\mathbf{I}$. Let ρ be an independent random variable taking the two positive real values a and b each with probability $1/2$. Let $\eta \in \{-1, +1\}$ be a Rademacher random variable independent from U and ρ.

$$\text{Let } X = \rho U \text{ and } Y = \langle \theta_*, X\rangle + \eta\|X\|^{-1} = \langle \theta_*, X\rangle + \eta/\rho.$$

In this case

$$G = \mathbb{E}\big(XX^\top\big) = \mathbb{E}\big(\rho^2\big)\mathbb{E}\big(UU^\top\big) = \frac{a^2+b^2}{2d}\mathbf{I},$$
$$C = \frac{2d}{a^2+b^2},$$
$$\text{and } R(\theta_*) = \frac{a^{-2}+b^{-2}}{2},$$
$$\text{so that } \frac{C}{dR(\theta_*)} = \frac{4}{(a^2+b^2)(a^{-2}+b^{-2})} = \frac{4}{2 + a^2/b^2 + b^2/a^2}$$

can take any value in the range $]0, 1]$ depending on the value of the ratio a/b. Meanwhile, it is easy to check that the hypotheses of Proposition 4.12 are fulfilled for any p and $q \in [1, 2]$ so that equation (23) on page 40 is satisfied for any $q \in]1, 2[$.

A. Proof of Proposition 1.1 on page 7

Let us assume without loss of generality that G is of full rank. Indeed it is easy to see that $X \in \mathbf{Im}(G)$ almost surely. This comes from the fact that for any $\theta \in \mathbf{Ker}(G)$, $\mathbb{E}\big(\langle \theta, X\rangle^2\big) = 0$, and therefore that $\mathbb{P}\big(\langle \theta, X\rangle = 0\big) = 1$.

Taking a basis of $\mathbf{Ker}(G)$, we obtain that $\mathbb{P}\big(X \in \mathbf{Ker}(G)^\perp\big) = 1$. Remark now that $\mathbf{Ker}(G)^\perp = \mathbf{Im}(G)$, since G is symmetric, so that $\mathbb{P}\big(x \in \mathbf{Im}(G)\big) = 1$. Restricting the state space to $\mathbf{Im}(G)$ and considering the coordinates of X in some orthonormal basis of $\mathbf{Im}(G)$ sets us back to the case where $\mathbf{Ker}(G) = \{0\}$.

Let us consider the new variables $\widetilde{\theta} = G^{1/2}\theta$, and $\widetilde{X} = G^{-1/2}X$. Working with $(\widetilde{\theta}, \widetilde{X})$ instead of (θ, X) is the same as assuming that $G = \mathbf{I}$ (the identity matrix of rank d), and consequently that $N(\theta) = \|\theta\|^2$. This is what we will do in the following of this proof, keeping in mind that with this convention, although $\langle \theta, X_i \rangle$ is still observable, X_i and θ themselves are not.

For any $\theta \in \mathbb{R}^d$, let us consider $\pi_\theta = \mathcal{N}(\theta, \beta^{-1}\,\mathbf{I})$, the Gaussian distribution with mean θ and covariance matrix $\beta^{-1}\,\mathbf{I}$.

In order to apply some PAC-Bayesian inequality, we introduce into the computations the perturbation of θ defined by π_θ.

We are going to prove a succession of lemmas leading to

Proposition A.1 *Let us introduce the numerical constant*

$$c = \frac{15}{8\log(2)(\sqrt{2}-1)} \exp\left(\frac{1+2\sqrt{2}}{2}\right) \leq 44.3.$$

For any real parameter $\lambda > 0$, any $x, \theta \in \mathbb{R}^d$,

$$\psi\Big\{\lambda\big[\langle\theta, x\rangle^2 - 1\big]\Big\} \leq \int \log\Bigg\{1 + \lambda\left[\langle\theta', x\rangle^2 - 1 - \frac{\|x\|^2}{\beta}\right] + \frac{\lambda^2}{2}\left[\langle\theta', x\rangle^2 - 1 - \frac{\|x\|^2}{\beta}\right]^2 + \frac{c\lambda^2\|x\|^2}{\beta}\left(\langle\theta', x\rangle^2 + \frac{\|x\|^2}{2\beta}\right)\Bigg\} \mathrm{d}\pi_\theta(\theta').$$

Let us introduce the function

$$\chi(x) = \begin{cases} \psi(x), & x \leq x_1, \\ y_1 + p_1(x - x_1) - (x - x_1)^2/8, & x_1 \leq x \leq x_1 + 4p_1, \\ y_1 + 2p_1^2, & x \geq x_1 + 4p_1, \end{cases}$$

where $x_1 \in [0,1], y_1, p_1$ are defined by the conditions $\psi''(x_1) = -1/4$, $y_1 = \psi(x_1)$, and $p_1 = \psi'(x_1)$.

Since $\psi''(x)$ continues to decrease after x_1, whereas $\chi''(x)$ remains constant, until $\chi'(z) = 0$, afterwhich the function χ is constant, we see that $\psi(z) \leq \chi(z)$ for all $z \in \mathbb{R}$. On the other hand,

$$\chi(z) \leq \log\big(1 + z + z^2/2\big).$$

Indeed, we already saw that this is the case for ψ, and as the function $f(z) = \log\bigl(1 + z + z^2/2\bigr)$ is such that $f(x_1) \geq \chi(x_1) = \psi(x_1)$, $f'(x_1) \geq \chi'(x_1) = \psi'(x_1)$, and $\inf f'' = -1/4$, f is above χ on the right-hand side of x_1 also.

Starting from the expressions

$$\psi'(z) = \frac{1-z}{1-z+z^2/2} \quad \text{and} \quad \psi''(z) = \frac{-z+z^2/2}{1-z+z^2/2},$$

it is easy to compute x_1, y_1 and p_1. We obtain

$$\begin{aligned} x_1 &= 1 - \sqrt{4\sqrt{2}-5}, \\ y_1 &= -\log\bigl[2(\sqrt{2}-1)\bigr], \\ p_1 &= \frac{\sqrt{4\sqrt{2}-5}}{2(\sqrt{2}-1)}, \\ \sup_{z\in\mathbb{R}} \chi(z) &= y_1 + 2p_1^2 = \frac{1+2\sqrt{2}}{2} - \log\bigl[2(\sqrt{2}-1)\bigr]. \end{aligned}$$

We start with a bound comparing perturbations of the parameter θ inside and outside of χ.

Lemma A.2 *For any $\rho \in \mathcal{M}_+^1(\Theta)$ and any $h \in \mathbb{L}^1(\rho)$,*

$$\chi\biggl(\int h \,\mathrm{d}\rho\biggr) \leq \int \chi(h)\,\mathrm{d}\rho + \frac{1}{8}\,\mathbf{Var}\bigl(h\,\mathrm{d}\rho\bigr),$$

where by definition

$$\mathbf{Var}\bigl(h\,\mathrm{d}\rho\bigr) = \int \biggl(h - \int h\,\mathrm{d}\rho\biggr)^2 \mathrm{d}\rho(\theta) \in \mathbb{R}\cup+\infty.$$

Proof. Let us consider the function

$$g(y) = \chi(y) + \frac{1}{8}\biggl(y - \int h\;\mathrm{d}\rho\biggr)^2, \qquad y \in \mathbb{R}.$$

Since $\inf \chi'' = 1/4$, the function g is convex. Jensen's inequality shows that

$$g\biggl(\int h\,\mathrm{d}\rho\biggr) \leq \int g(h)\,\mathrm{d}\rho,$$

and we conclude, remarking that

$$g\biggl(\int h\,\mathrm{d}\rho\biggr) = \chi\biggl(\int h\,\mathrm{d}\rho\biggr).$$

□

LEMMA A.3 *For any $\rho \in \mathcal{M}_+^1(\Theta)$ and any $h \in \mathbb{L}^1(\rho)$,*

$$\psi\biggl(\int h \,\mathrm{d}\rho\biggr) \leq \int \chi(h) \,\mathrm{d}\rho + \min\Bigl\{\log(4), \frac{1}{8}\,\mathbf{Var}\bigl(h \,\mathrm{d}\rho\bigr)\Bigr\}.$$

PROOF.

$$\psi\biggl(\int h \,\mathrm{d}\rho\biggr) \leq \sup \psi + \int \chi(h) \,\mathrm{d}\rho - \inf \chi = \int \chi(h) \,\mathrm{d}\rho + \log(4).$$

We obtain this lemma by combining this inequality with the previous lemma (since $\psi(z) \leq \chi(z)$ for any $z \in \mathbb{R}$). $\square$

Applying this to our problem gives, for any $\theta \in \mathbb{R}^d$,

$$\begin{aligned}
\psi\Bigl\{\lambda\bigl[\langle \theta, x\rangle^2 - 1\bigr]\Bigr\} &= \psi\left\{\lambda\left[\int \langle \theta', x\rangle^2 \,\mathrm{d}\pi_\theta(\theta') - \frac{\|x\|^2}{\beta} - 1\right]\right\} \\
&\leq \int \chi\left\{\lambda\left[\langle \theta', x\rangle^2 - \frac{\|x\|^2}{\beta} - 1\right]\right\} \mathrm{d}\pi_\theta(\theta') \\
&\quad + \min\Bigl\{\log(4), \frac{\lambda^2}{8}\,\mathbf{Var}\bigl[\langle x, \theta'\rangle^2 \,\mathrm{d}\pi_\theta(\theta')\bigr]\Bigr\} \\
&= \int \chi\left\{\lambda\left[\langle \theta', x\rangle^2 - \frac{\|x\|^2}{\beta} - 1\right]\right\} \mathrm{d}\pi_\theta(\theta') \\
&\qquad + \min\left\{\log(4), \frac{\lambda^2\|x\|^2\langle \theta, x\rangle^2}{2\beta} + \frac{\lambda^2\|x\|^4}{4\beta^2}\right\},
\end{aligned}$$

where we have used the fact that when $W \sim \mathcal{N}(0, \sigma^2)$,

$$\begin{aligned}
\mathbf{Var}\bigl[(m+W)^2\bigr] &= \mathbf{Var}\bigl(W^2 + 2mW\bigr) \\
&= \mathbb{E}\Bigl[\bigl(W^2 + 2mW\bigr)^2\Bigr] - \mathbb{E}\bigl(W^2 + 2mW\bigr)^2 \\
&= \mathbb{E}\bigl(W^4\bigr) + 4m^2\sigma^2 - \sigma^4 = 2\sigma^4 + 4m^2\sigma^2,
\end{aligned}$$

to compute $\mathbf{Var}\bigl[\langle x, \theta'\rangle^2 \,\mathrm{d}\pi_\theta(\theta')\bigr]$.
Remark now that for any positive real numbers a, b, and c,

$$\min\bigl\{a, bm^2 + c\bigr\} \leq \min\bigl\{a, b(m+W)^2 + c\bigr\} + \min\bigl\{a, b(m-W)^2 + c\bigr\},$$

so that, the distribution of W and $-W$ being the same,

$$\min\bigl\{a, bm^2 + c\bigr\} \leq 2\mathbb{E}\Bigl[\min\bigl\{a, b(m+W)^2 + c\bigr\}\Bigr].$$

Applying this to the gaussian distribution π_θ, we obtain that

$$\psi\big\{\lambda\big[\langle\theta,x\rangle^2-1\big]\big\}\leq\int\chi\bigg\{\lambda\bigg[\langle\theta',x\rangle^2-\frac{\|x\|^2}{\beta}-1\bigg]\bigg\}\,\mathrm{d}\pi_\theta(\theta')\\+\int\min\bigg\{4\log(2),\frac{\lambda^2\|x\|^2\langle\theta',x\rangle^2}{\beta}+\frac{\lambda^2\|x\|^4}{2\beta^2}\bigg\}\,\mathrm{d}\pi_\theta(\theta').$$

We are now going to use the fact that for any $a,b,y\in\mathbb{R}_+$, such that $y\leq b$,

$$\log(a)+\min\{b,y\}=\log\Big[a\exp\big(\min\{b,y\}\big)\Big]\\\leq\log\bigg\{a+\min\{b,y\}\frac{a\big(\exp(b)-1\big)}{b}\bigg\}\leq\log\bigg\{a+y\frac{a\big(\exp(b)-1\big)}{b}\bigg\}.$$

Applying this inequality to $a=\exp\big(\chi(z)\big)$ and reminding that $\exp\big[\chi(z)\big]\leq 1+z+z^2$, we obtain that

$$\chi(z)+\min\{b,y\}\leq\log\bigg\{1+z+z^2/2+y\frac{\exp\big(\sup\chi\big)\big(\exp(b)-1\big)}{b}\bigg\}.$$

Coming back to our problem, where we can take $b=4\log(2)$, we get

$$\psi\Big\{\lambda\big[\langle\theta,x\rangle^2-1\big]\Big\}\\\leq\int\log\bigg\{1+\lambda\bigg[\langle\theta',x\rangle^2-1-\frac{\|x\|^2}{\beta}\bigg]+\frac{\lambda^2}{2}\bigg[\langle\theta',x\rangle^2-1-\frac{\|x\|^2}{\beta}\bigg]^2\\+\frac{c\lambda^2\|x\|^2}{\beta}\bigg[\langle\theta',x\rangle^2+\frac{\|x\|^2}{2\beta}\bigg]\bigg\}\,\mathrm{d}\pi_\theta(\theta'),$$

where

$$c=\frac{15}{8\log(2)(\sqrt{2}-1)}\exp\bigg(\frac{1+2\sqrt{2}}{2}\bigg),$$

as announced in Proposition A.1 on page 43.

Now that we have compared $\psi\Big\{\lambda\big[\langle\theta,x\rangle^2-1\big]\Big\}$ with an expectation with respect to a Gaussian perturbation of the parameter θ, we are prepared to use the following PAC-Bayes inequality.

LEMMA A.4 *Let us consider two measurable spaces $\mathcal{X}$ and Θ and an i.i.d. sample $(X_i)_{i=1}^n\in\mathcal{X}^n$ with probability distribution $\mathbb{P}^{\otimes n}$. Let $\overline{\mathbb{P}}=\frac{1}{n}\sum_{i=1}^n\delta_{X_i}$ be its empirical measure. Let us consider a prior probability measure $\nu\in$*

$\mathcal{M}_+^1(\Theta)$. *For any measurable function* $f : \mathcal{X} \times \Theta \to [a, +\infty[$, *where* $-1 < a < +\infty$, *with probability at least* $1 - \epsilon$ *according to the sample distribution* $\mathbb{P}^{\otimes n}$, *for any posterior probability measure* $\rho \in \mathcal{M}_+^1(\Theta)$ *such that* $\mathcal{K}(\rho, \nu) < +\infty$,

$$\int \log\bigl[1 + f(x,\theta)\bigr] \,\mathrm{d}\rho(\theta) \,\mathrm{d}\overline{\mathbb{P}}(x) \leq \int f(x,\theta) \,\mathrm{d}\mathbb{P}(x) \,\mathrm{d}\rho(\theta) + \frac{\mathcal{K}(\rho,\nu) - \log(\epsilon)}{n}. \quad (24)$$

Let us remark that in this lemma and in the following of this paper we will encounter events that are the union or the intersection of an uncountable family of measurable sets, as well as suprema of uncountable families of functions. To give a meaning to this, consider that when we write $\mathbb{P}(A) \geq 1 - \epsilon$, we mean that there exists a measurable set $B \subset A$ such that $\mathbb{P}(B) \geq 1 - \epsilon$. In the same way when we write

$$\int h \,\mathrm{d}\mathbb{P} \leq \eta,$$

we mean that there is a measurale function g such that $h(x) \leq g(x)$ for all $x \in \mathcal{X}$ and

$$\int g \,\mathrm{d}\mathbb{P} \leq \eta.$$

The proof of Lemma A.4 on the previous page requires a succession of preliminary results.

Lemma A.5 *For any measurable space* Θ, *any upper bounded measurable function* $h : \Theta \to \mathbb{R}$, *any probability measures* $\rho, \pi \in \mathcal{M}_+^1(\Theta)$,

$$\int h \,\mathrm{d}\rho - \mathcal{K}(\rho,\pi) \leq \log\biggl(\int \exp(h) \,\mathrm{d}\pi\biggr).$$

Proof. Let us introduce the probability measure $\pi_{\exp(h)} \ll \pi$ with density $\dfrac{\mathrm{d}\pi_{\exp(h)}}{\mathrm{d}\pi} = \dfrac{\exp(h)}{\int \exp(h) \,\mathrm{d}\pi}$. We can check that

$$\int h \,\mathrm{d}\rho - \mathcal{K}(\rho,\pi) - \log\biggl(\int \exp(h) \,\mathrm{d}\pi\biggr) = -\mathcal{K}(\rho, \pi_{\exp(h)}) \leq 0,$$

(where it is possible that $\mathcal{K}(\rho,\pi) = +\infty$, $\mathcal{K}(\rho, \pi_{\exp(h)}) = +\infty$, or $\int h \,\mathrm{d}\rho = -\infty$). □

LEMMA A.6 *Let us consider some measurable state space $\mathcal{X}$, some measurable parameter space Θ, a prior probability measure ν. For any measurable function $g : \mathcal{X} \times \Theta \to \mathbb{R}$, such that*

$$\int \exp\bigl[g(x,\theta)\bigr]\,\mathrm{d}\mathbb{P}(x) \leq 1, \qquad \theta \in \Theta,$$

$$\int \exp\biggl(\sup_{\substack{\rho \in \mathcal{M}_+^1(\Theta),\\ \mathcal{K}(\rho,\nu)<+\infty}} n \int g(x,\theta)\,\mathrm{d}\overline{\mathbb{P}}(x)\,\mathrm{d}\rho(\theta) - \mathcal{K}(\rho,\nu)\biggr)\,\mathrm{d}\mathbb{P}^{\otimes n} \leq 1$$

Therefore, with probability at least $1-\epsilon$, for any $\rho \in \mathcal{M}_+^1(\Theta)$ such that $\mathcal{K}(\rho,\nu) < +\infty$,

$$\int g(x,\theta)\,\mathrm{d}\overline{\mathbb{P}}(x)\,\mathrm{d}\rho(\theta) < \frac{\mathcal{K}(\rho,\nu) - \log(\epsilon)}{n}.$$

Let us give some needed explanations about the meaning of this lemma. The assumptions do not imply that $\theta \mapsto g(x,\theta) \in \mathbb{L}^1(\rho)$. As we will see in the proof, the lemma is true without further assumptions if we use the convention

$$\int g(x,\theta)\,\mathrm{d}\rho(\theta) = -\infty \text{ when } \int \min\bigl\{g(x,\theta),0\bigr\}\,\mathrm{d}\rho(\theta) = -\infty,$$

and is defined otherwise as usual (so that the integral is equal to $+\infty$ when the integral of the positive part of the integrand is infinite and the integral of its negative part is finite).

PROOF. Let us put $g_k(x,\theta) = \min\{g(x,\theta),k\}$, $k \in \mathbb{N}$. Applying the monotone convergence theorem, the previous lemma, and the monotone convergence theorem again, we see that

$$\begin{aligned}
&\exp\biggl(\sup_{\substack{\rho \in \mathcal{M}_+^1(\Theta),\\ \mathcal{K}(\rho,\nu)<+\infty}} n \int g(x,\theta)\,\mathrm{d}\overline{\mathbb{P}}(x)\,\mathrm{d}\rho(\theta) - \mathcal{K}(\rho,\nu)\biggr)\\
&\qquad = \sup_{k\in\mathbb{N}} \exp\biggl(\sup_{\substack{\rho \in \mathcal{M}_+^1(\Theta),\\ \mathcal{K}(\rho,\nu)<+\infty}} n \int g_k(x,\theta)\,\mathrm{d}\overline{\mathbb{P}}(x)\,\mathrm{d}\rho(\theta) - \mathcal{K}(\rho,\nu)\biggr)\\
&\qquad\qquad \leq \sup_{k\in\mathbb{N}} \biggl[\int \exp\biggl(n \int g_k(x,\theta)\,\mathrm{d}\overline{\mathbb{P}}(x)\biggr)\,\mathrm{d}\nu(\theta)\biggr]\\
&\qquad\qquad\qquad = \int \exp\biggl(n \int g(x,\theta)\,\mathrm{d}\overline{\mathbb{P}}(x)\biggr)\,\mathrm{d}\nu(\theta).
\end{aligned}$$

According to Fubini's theorem, the right-hand side of the last equality is a measurable function of the sample. Thus, with our conventions about the use of integrals of non-measurable functions, we can write

$$\begin{aligned}\int \exp\Bigg(\sup_{\substack{\rho\in\mathcal{M}_+^1(\Theta),\\ \mathcal{K}(\rho,\nu)<+\infty}} & n\int g(x,\theta)\,\mathrm{d}\overline{\mathbb{P}}(x)\,\mathrm{d}\rho(\theta) - \mathcal{K}(\rho,\nu)\Bigg)\,\mathrm{d}\mathbb{P}^{\otimes n}\\ &\leq \int \exp\left(n\int g(x,\theta)\,\mathrm{d}\overline{\mathbb{P}}(x)\right)\mathrm{d}\nu(\theta)\,\mathrm{d}\mathbb{P}^{\otimes n}\\ &= \int \exp\left(n\int g(x,\theta)\,\mathrm{d}\overline{\mathbb{P}}(x)\right)\mathrm{d}\mathbb{P}^{\otimes n}\,\mathrm{d}\nu(\theta)\\ &= \int\left(\int \exp\bigl[g(x,\theta)\bigr]\,\mathrm{d}\mathbb{P}(x)\right)^n \mathrm{d}\nu(\theta) \leq 1,\end{aligned}$$

where we have used Fubini's theorem for non-negative functions again and the independence of the sample.

This proves the first inequality of the lemma.

To prove the second one, let us consider the two events

$$\begin{aligned}A = \Bigg\{ &\int g(x,\theta)\,\mathrm{d}\overline{\mathbb{P}}(x)\,\mathrm{d}\rho(\theta) < \frac{\mathcal{K}(\rho,\nu)-\log(\epsilon)}{n},\\ &\qquad\qquad \rho\in\mathcal{M}_+^1(\Theta), \mathcal{K}(\rho,\nu)<+\infty\Bigg\},\\ B = \Bigg\{ &\log\left[\int \exp\left(n\int g(x,\theta)\,\mathrm{d}\overline{\mathbb{P}}(x)\right)\mathrm{d}\nu(\theta)\right] < -\log(\epsilon)\Bigg\}.\end{aligned}$$

According to the first part of the proof, $B\subset A$. Moreover B is measurable. Let us introduce the random variable

$$W = \log\left[\int \exp\left(n\int g(x,\theta)\,\mathrm{d}\overline{\mathbb{P}}\right)\mathrm{d}\nu(\theta)\right] + \log(\epsilon).$$

It is measurable and $B=\{W<0\}$, so that $\mathbb{P}^{\otimes n}(B) = 1-\mathbb{P}^{\otimes n}\bigl(W\geq 0\bigr)$. We end the proof using the exponential Markov inequality

$$\mathbb{P}^{\otimes n}\bigl(W\geq 0\bigr) \leq \int \exp(W)\,\mathrm{d}\mathbb{P}^{\otimes n} \leq \epsilon,$$

the last inequality being a consequence of the first part of the proof. □

Proof of Lemma A.4 on page 46

Let us put $m(\theta) = \int f(x,\theta)\,\mathrm{d}\mathbb{P}(x) \in [a,+\infty]$. Let us consider

$$g(x,\theta) = \mathbb{1}\big[m(\theta) < +\infty\big]\Big[\log\big(1+f(x,\theta)\big) - m(\theta)\Big]$$

We can check that

$$\int \exp\big[g(x,\theta)\big]\,\mathrm{d}\mathbb{P}(x) = \begin{cases} \dfrac{1+m(\theta)}{\exp\big[m(\theta)\big]}, & \text{when } m(\theta) < +\infty \\ 1, & \text{otherwise.} \end{cases}$$

Since $1+m \leq \exp(m)$, from the convexity of the exponential function, we see that

$$\int \exp\big[g(x,\theta)\big]\,\mathrm{d}\mathbb{P}(x) \leq 1.$$

Therefore, we can apply the previous lemma to this choice of g. It shows that with probability at least $1-\epsilon$, for any $\rho \in \mathcal{M}_+^1(\Theta)$, such that $\mathcal{K}(\rho,\nu) < +\infty$,

$$\int \mathbb{1}\big[m(\theta') < +\infty\big]\Big\{\log\big[1+f(x,\theta')\big] - m(\theta')\Big\}\,\mathrm{d}\overline{\mathbb{P}}(x)\,\mathrm{d}\rho_\theta(\theta') < B(\rho,\epsilon), \quad (25)$$

where $B(\rho,\epsilon) = \dfrac{\mathcal{K}(\rho,\nu) - \log(\epsilon)}{n}$. We can then remark that in the case when $\int m(\theta)\,\mathrm{d}\rho(\theta) < +\infty$, $\int \mathbb{1}\big[m(\theta) = +\infty\big]\,\mathrm{d}\rho(\theta) = 0$ and the left-hand side of equation (25) is equal to

$$\int \log\big[1+f(x,\theta)\big]\,\mathrm{d}\overline{\mathbb{P}}(x)\,\mathrm{d}\rho(\theta) - \int m(\theta)\,\mathrm{d}\rho(\theta),$$

so that in this case equation (25) is equation (24) on page 47 with a strict inequality, and therefore implies equation (24). On the other hand, when $\int m(\theta)\,\mathrm{d}\rho(\theta) = +\infty$, inequality (24) on page 47 is also true because its right-hand side is equal to $+\infty$. □

Let us now apply Lemma A.4 on page 46 to our problem, choosing $\mathcal{X} = \Theta = \mathbb{R}^d$, $\rho = \pi_\theta$ and $\nu = \pi_0$. It proves that, with probability at least $1-\epsilon$, for any $\theta \in \mathbb{R}^d$,

$$\int \psi\Big\{\lambda\big[\langle\theta,x\rangle^2 - 1\big]\Big\}\,\mathrm{d}\overline{\mathbb{P}}(x) \leq \int\Bigg\{\lambda\bigg(\langle\theta',x\rangle^2 - 1 - \frac{\|x\|^2}{\beta}\bigg) + \frac{\lambda^2}{2}\bigg(\langle\theta',x\rangle^2 - 1 - \frac{\|x\|^2}{\beta}\bigg)^2 + \frac{c\lambda^2\|x\|^2}{\beta}\bigg(\langle\theta',x\rangle^2 + \frac{\|x\|^2}{2\beta}\bigg)\Bigg\}\,\mathrm{d}\pi_\theta(\theta')\,\mathrm{d}\mathbb{P}(x)$$

$$+ \frac{\mathcal{K}(\pi_\theta, \pi_0) - \log(\epsilon)}{n}.$$

We can now compute $\mathcal{K}(\pi_\theta, \pi_0) = \dfrac{\beta\|\theta\|^2}{2} = \displaystyle\int \frac{\beta\langle\theta, x\rangle^2}{2}\,\mathrm{d}\mathbb{P}(x)$, taking here into account the change of variables that turns the Gram matrix into the identity. Integrating then explicitly with respect to the Gaussian measure π_θ, we get with probability at least $1-\epsilon$, for all $\theta \in \mathbb{R}^d$, that

$$\begin{aligned}\int \psi\Big\{\lambda\big[\langle\theta, x\rangle^2 - 1\big]\Big\}\,\mathrm{d}\overline{\mathbb{P}}(x) \le \int \Bigg\{ & \lambda\big(\langle\theta, x\rangle^2 - 1\big) \\ & + \frac{\lambda^2}{2}\left[\big(\langle\theta, x\rangle^2 - 1\big)^2 + \frac{4\|x\|^2}{\beta}\langle\theta, x\rangle^2 + \frac{2\|x\|^4}{\beta^2}\right] \\ & + \frac{c\lambda^2\|x\|^2}{\beta}\left(\langle\theta, x\rangle^2 + \frac{3\|x\|^2}{2\beta}\right) + \frac{\beta\langle\theta, x\rangle^2}{2n}\Bigg\}\,\mathrm{d}\mathbb{P}(x) - \frac{\log(\epsilon)}{n}.\end{aligned}$$

Let us introduce

$$\begin{aligned} s_4^4 &= \int \|x\|^4\,\mathrm{d}\mathbb{P}(x), \\ \kappa &= \sup\left\{\int \langle\theta, x\rangle^4\,\mathrm{d}\mathbb{P}(x)\,;\ \theta \in \mathbb{R}^d,\ \int \langle\theta, x\rangle^2\,\mathrm{d}\mathbb{P}(x) \le 1\right\}. \end{aligned}$$

LEMMA A.7 *The coefficients s_4 and κ are such that*

$$s_4^4 \le \kappa\left(\int \|x\|^2\,\mathrm{d}\mathbb{P}(x)\right)^2 = \kappa d^2.$$

PROOF. Taking coordinates, we can write

$$\begin{aligned} s_4^4 = \int\left(\sum_{i=1}^d x_i^2\right)^2\mathrm{d}\mathbb{P}(x) &= \sum_{1\le i,j\le d}\int x_i^2 x_j^2\,\mathrm{d}\mathbb{P}(x) \\ &\le \sum_{1\le i,j\le d}\left(\int x_i^4\,\mathrm{d}\mathbb{P}(x)\right)^{1/2}\left(\int x_j^4\,\mathrm{d}\mathbb{P}(x)\right)^{1/2} \\ &\le \kappa\sum_{1\le i,j\le d}\left(\int x_i^2\,\mathrm{d}\mathbb{P}(x)\right)\left(\int x_j^2\,\mathrm{d}\mathbb{P}(x)\right) \\ &= \kappa\left(\int \|x\|^2\,\mathrm{d}\mathbb{P}(x)\right)^2. \end{aligned}$$

□

Using the Cauchy-Schwarz inequality, we obtain with probability at least $1-\epsilon$, for any $\theta \in \mathbb{R}$,

$$r_\lambda(\theta) \le \left[1 + (\kappa-1)\lambda + \frac{(2+c)\lambda s_4^2\sqrt{\kappa}}{\beta} + \frac{\beta}{2n\lambda}\right]\big[N(\theta)-1\big] + \frac{\kappa\lambda}{2}\big[N(\theta)-1\big]^2 + \frac{(\kappa-1)\lambda}{2} + \frac{(2+c)\sqrt{\kappa}\lambda s_4^2}{\beta} + \frac{(2+3c)\lambda s_4^4}{2\beta^2} + \frac{\beta}{2n\lambda} - \frac{\log(\epsilon)}{n\lambda}.$$

Let us choose some special values for β, the strength of the perturbation of the parameter, and λ, the scale parameter of the influence function ψ. Let us also introduce some more concise notations. Let us put accordingly

$$\begin{aligned}
\beta &= \lambda s_4 \kappa^{1/4}\sqrt{2(2+c)n},\\
\lambda &= \sqrt{\frac{2}{(\kappa-1)n}\left[\log\big(\epsilon^{-1}\big) + \frac{(2+3c)s_4^2}{4\sqrt{\kappa}(2+c)}\right]},\\
\eta &= (\kappa-1)\lambda = \sqrt{\frac{2(\kappa-1)}{n}\left[\log\big(\epsilon^{-1}\big) + \frac{(2+3c)s_4^2}{4\sqrt{\kappa}(2+c)}\right]},\\
\gamma &= \sqrt{\frac{2(2+c)s_4^2\sqrt{\kappa}}{n}},\\
\mu &= \eta + \gamma,\\
\xi &= \frac{\kappa\eta}{2(\kappa-1)}.
\end{aligned}$$

Let us remark that, according to Lemma A.7 on page 51, $s_4^2 \le \sqrt{\kappa} d$, so that

$$\mu \le \sqrt{\frac{2(\kappa-1)}{n}\left[\log(\epsilon^{-1}) + \frac{(2+3c)d}{4(2+c)}\right]} + \sqrt{\frac{2(2+c)\kappa d}{n}}$$

Numerically,

$$\mu \le \sqrt{\frac{2(\kappa-1)}{n}\big[\log(\epsilon^{-1}) + 0.73\, d\big]} + 6.81\sqrt{\frac{2\kappa\, d}{n}}$$

Proposition A.8 *With probability at least $1-\epsilon$, for any $\theta \in \mathbb{R}^d$,*

$$r_\lambda(\theta) \le \xi\big[N(\theta)-1\big]^2 + (1+\mu)\big[N(\theta)-1\big] + \mu.$$

Let us now study the reverse inequality. Starting from

$$\psi\Big\{\lambda\big[1-\langle\theta,x\rangle^2\big]\Big\} = \psi\left\{\lambda\left[1+\frac{\|x\|^2}{\beta}-\int\langle\theta',x\rangle^2\,\mathrm{d}\pi_\theta(\theta')\right]\right\}$$
$$\leq \int\chi\left\{\lambda\left[1+\frac{\|x\|^2}{\beta}-\langle\theta',x\rangle^2\right]\right\}\mathrm{d}\pi_\theta(\theta')$$
$$+\min\Big\{\log(4),\frac{\lambda^2}{8}\,\mathbf{Var}\big[\langle\theta',x\rangle^2\,\mathrm{d}\pi_\theta(\theta')\big]\Big\},$$

and proceeding in the same way as previously, we obtain

PROPOSITION A.9 *With probability at least $1-\epsilon$, for any $\theta\in\mathbb{R}^d$,*

$$r_\lambda(\theta)\geq -\xi\big[N(\theta)-1\big]^2+(1-\mu)\big[N(\theta)-1\big]-\mu.$$

PROPOSITION A.10 *Assume that*

$$2\mu+\xi<1, \tag{26}$$

which can be written as

$$n>\left\{\sqrt{8(2+c)}\kappa^{1/4}s_4 + \left(\frac{5}{2}+\frac{1}{2(\kappa-1)}\right)\sqrt{2(\kappa-1)\left[\log(\epsilon^{-1})+\frac{(2+3c)s_4^2}{4\sqrt{\kappa}(2+c)}\right]}\right\}^2.$$

A sufficient numerical condition implying the one above is that

$$n>\left[20\sqrt{\kappa d}+\left(\frac{5}{2}+\frac{1}{2(\kappa-1)}\right)\sqrt{2(\kappa-1)\big[\log(\epsilon^{-1})+0.73\,d\big]}\right]^2.$$

With probability at least $1-2\epsilon$, for any $\theta\in\mathbb{R}^d$,

$$\left(1-\frac{\mu}{1-2\mu}\right)\widehat{N}(\theta)\leq N(\theta)\leq\left(1+\frac{\mu}{1-2\mu}\right)\widehat{N}(\theta),$$

$$\textit{where}\quad c=\frac{15}{8\log(2)\big(\sqrt{2}-1\big)}\exp\left(\frac{1+2\sqrt{2}}{2}\right),$$

$$\text{and } \mu = \sqrt{\frac{2(\kappa-1)}{n}\left[\log\bigl(\epsilon^{-1}\bigr) + \frac{(2+3c)s_4^2}{4\sqrt{\kappa}(2+c)}\right]} + \sqrt{\frac{2(2+c)\sqrt{\kappa}s_4^2}{n}}$$
$$\leq \sqrt{\frac{2(\kappa-1)}{n}\left[\log\bigl(\epsilon^{-1}\bigr) + \frac{(2+3c)d}{4(2+c)}\right]} + 2\sqrt{\frac{2(2+c)\kappa d}{n}}$$
$$\leq \sqrt{\frac{2(\kappa-1)}{n}\bigl[\log\bigl(\epsilon^{-1}\bigr) + 0.73d\bigr]} + 6.81\sqrt{\frac{2\kappa d}{n}}.$$

Proof. Assume that both inequalities of Proposition A.8 and A.9 hold for any $\theta \in \mathbb{R}^d$, which happens at least with probability $1-2\epsilon$. In this case, when $N(\theta) = 0$, for any $\alpha \in \mathbb{R}_+$,

$$r_\lambda(\alpha\theta) \leq -1 + \xi < 0,$$

so that $\widehat{\alpha} = +\infty$, and $\widehat{N}(\theta) = 0 = N(\theta)$. Assume now that $N(\theta) \neq 0$. In this case, there is $\alpha = \sqrt{\dfrac{2}{N(\theta)}}$ such that $N(\alpha\theta) = 2$ and therefore such that

$$r_\lambda(\alpha\theta) \geq 1 - 2\mu - \xi > 0,$$

proving that $\widehat{\alpha}(\theta) < +\infty$, and consequently that $r\bigl[\widehat{\alpha}(\theta)\theta\bigr] = 0$, as explained before. For any $\alpha \in I = [0, \widehat{\alpha}(\theta)]$, $r_\lambda(\alpha\theta) \leq 0$, because $\alpha \mapsto r_\lambda(\alpha\theta)$ is non-decreasing. Thus $N(\alpha\theta) - 1$ is solution of the following quadratic inequality in the unknown variable z :

$$0 \geq -\xi z^2 + (1-\mu)z - \mu. \tag{27}$$

Since the right-hand side of this inequality is positive when $z = 1$, from assumption (26) on page 53, its discriminant

$$\Delta = (1-\mu)^2 - 4\xi\mu$$

is strictly positive. Let us consider the interval

$$J = \frac{1}{2\xi}\Bigl(1 - \mu + \sqrt{\Delta}\,\Bigr]-1, +1\Bigl[\Bigr).$$

We obtain that

$$\bigl[N(\alpha\theta) - 1\bigr] \notin J, \qquad \alpha \in [0, \widehat{\alpha}(\theta)].$$

Since $\alpha \mapsto N(\alpha\theta) - 1$ is continuous and $-1 \leq \inf J$, necessarily

$$N\bigl[\widehat{\alpha}(\theta)\theta\bigr] - 1 \leq \inf J \leq \frac{\mu}{1-2\mu}.$$

Indeed, when $z = \dfrac{\mu}{1-2\mu}$, the right-hand side of inequality (27) is equal to

$$\frac{1-2\mu-\xi}{(1-2\mu)^2} > 0$$

according to assumption (26) on page 53, so that this value of z belongs to J.

On the other hand, $1 - N\big[\widehat{\alpha}(\theta)\theta\big] \leq 1$ is also solution of

$$0 \leq \xi z^2 - (1+\mu)z + \mu.$$

When $z = 1$, the right-hand side of this inequality is strictly negative, due to assumption (26) on page 53, showing that the corresponding equality has two roots and that $1 - N\big[\widehat{\alpha}(\theta)\theta\big]$ is lower than its lowest root which is in turn lower than $\dfrac{\mu}{1-2\mu}$. Thus, as $N\big[\widehat{\alpha}(\theta)\theta\big] = \widehat{\alpha}(\theta)^2 N(\theta) = N(\theta)/\widehat{N}(\theta)$, we have proved that

$$\left|\frac{N(\theta)}{\widehat{N}(\theta)} - 1\right| \leq \frac{\mu}{1-2\mu},$$

from which Proposition A.10 on page 53, and therefore Proposition 1.1 on page 7, follow. □

B. Proof of Proposition 2.2 on page 11

Let Ω be the event of probability at least $1-2\epsilon$ that appears in Proposition 1.1 on page 7. Remark that on Ω,

$$\begin{aligned}
\int \Big(\langle\theta,x\rangle^2 \widehat{N}(\theta)^{-1} - 1\Big)_+^3 \,\mathrm{d}\overline{\mathbb{P}}(x) &\leq \int \frac{\langle\theta,x\rangle^6}{\widehat{N}(\theta)^3}\,\mathrm{d}\overline{\mathbb{P}}(x) \\
&\leq \frac{\max\limits_{i=1,\dots,n}\langle\theta,X_i\rangle^4}{\widehat{N}(\theta)^2}\int \frac{\langle\theta,x\rangle^2}{\widehat{N}(\theta)}\,\mathrm{d}\overline{\mathbb{P}}(x) \\
&\leq \frac{\max\limits_{i=1,\dots,n}\|G^{1/2}\theta\|^4\|G^{-1/2}X_i\|^4}{\widehat{N}(\theta)^2} \times \frac{\overline{N}(\theta)}{\widehat{N}(\theta)} \\
&= \frac{N(\theta)^2}{\widehat{N}(\theta)^2} \times R^4 \times \frac{\overline{N}(\theta)}{\widehat{N}(\theta)} \leq \big(1+\widehat{\delta}\big)^2 R^4 \frac{\overline{N}(\theta)}{\widehat{N}(\theta)}. \quad (28)
\end{aligned}$$

Remark also that the conclusions of Proposition 2.1 on page 10 hold on Ω, so that

$$-\overline{\delta}_-(\theta) \leq \frac{\overline{N}(\theta)}{\widehat{N}(\theta)} - 1 \leq \overline{\delta}_+(\theta).$$

From equation (28),

$$\overline{\delta}_+(\theta) \le \gamma_+ \frac{\overline{N}(\theta)}{\widehat{N}(\theta)},$$

so that

$$\frac{\overline{N}(\theta)}{\widehat{N}(\theta)} \le \frac{1}{(1-\gamma_+)_+}$$

and

$$\overline{\delta}_+(\theta) \le \frac{\gamma_+}{(1-\gamma_+)_+}.$$

This gives

$$-\gamma_- \le -\overline{\delta}_-(\theta) \le \frac{\overline{N}(\theta)}{\widehat{N}(\theta)} - 1 \le \overline{\delta}_+(\theta) \le \frac{\gamma_+}{(1-\gamma_+)_+}.$$

Using Proposition 1.1 again, we get that on the event Ω,

$$\frac{(1-\widehat{\delta})\overline{N}(\theta)}{1+\overline{\delta}_+(\theta)} \le (1-\widehat{\delta})\widehat{N}(\theta) \le N(\theta) \le (1+\widehat{\delta})\widehat{N}(\theta) \le \frac{(1+\widehat{\delta})\overline{N}(\theta)}{1-\overline{\delta}_-(\theta)}.$$

As a consequence

$$\begin{aligned} -\frac{\widehat{\delta}+\gamma_-}{1+\widehat{\delta}} &\le -\frac{\widehat{\delta}+\overline{\delta}_-(\theta)}{1+\widehat{\delta}} \le \frac{\overline{N}(\theta)}{N(\theta)} - 1 \le \frac{\widehat{\delta}+\overline{\delta}_+(\theta)}{1-\widehat{\delta}} \\ &\le \frac{\widehat{\delta}+\gamma_+/(1-\gamma_+)_+}{1-\widehat{\delta}} = \frac{1}{(1-\widehat{\delta})(1-\gamma_+)_+} - 1 \le \frac{\widehat{\delta}+\gamma_+}{(1-\widehat{\delta})(1-\gamma_+)_+}. \end{aligned}$$

C. Proof of Lemma 2.5 on page 16

Consider $u \in \mathbb{R}_+$ that we will choose appropriately and write the decomposition

$$\frac{1}{n}\sum_{i=1}^n W_i = \frac{1}{n}\sum_{i=1}^n \min\{W_i, u\} + \frac{1}{n}\sum_{i=1}^n \big(W_i - u\big)_+.$$

Using Bienaymé Chebyshev's inequality on one hand and Markov's inequality on the other hand, we obtain that with probability at least $1-2\epsilon$,

$$\frac{1}{n}\sum_{i=1}^n W_i \le \mathbb{E}(W) + \sqrt{\frac{\mathbb{E}\big(\min\{W,u\}^2\big)}{n\epsilon}} + \frac{\mathbb{E}\big[\big(W-u\big)_+\big]}{\epsilon}.$$

Let us now remark that

$$\min\{W,u\}^2 \leq u^{2-q}W^q \text{ and } \big(W-u\big)_+ \leq \frac{(q-1)^{q-1}}{q^q}u^{-(q-1)}W^q.$$

The first inequality is obvious. The second one can be checked remarking that the two functions of W meet at $W = u/(1-q^{-1})$, where they have the same derivative and that the function on the right-hand side being convex is necessarily above its tangent $W \mapsto W-u$, and therefore also above $\big(W-u\big)_+$, since it is positive. More precisely, if we compute the solution of the contact equations

$$W_c - u = aW_c^q \text{ and } 1 = qaW_c^{q-1},$$

we find

$$W_c = \frac{qu}{q-1} \text{ and } a = \frac{1}{q}\left(\frac{q-1}{qu}\right)^{q-1} = \frac{(q-1)^{q-1}}{q^q}u^{-(q-1)},$$

as claimed. Thus with probability at least $1-2\epsilon$,

$$\frac{1}{n}\sum_{i=1}^n W_i \leq \mathbb{E}(W) + \sqrt{\frac{u^{2-q}\mathbb{E}(W^q)}{n\epsilon}} + \frac{(q-1)^{q-1}\mathbb{E}\big(W^q\big)}{q^q u^{q-1}\epsilon}.$$

The optimal value of the threshold u can be found computing the derivative of the bound. It is

$$u = \frac{(q-1)^2\mathbb{E}\big(W^q\big)^{1/q}n^{1/q}}{q^2(1-q/2)^{2/q}\epsilon^{1/q}}.$$

Replacing u by its value gives as stated in the lemma

$$\frac{1}{n}\sum_{i=1}^n W_i \leq \mathbb{E}(W) + \frac{q^{q-1}\mathbb{E}(W^q)^{1/q}}{2(q-1)^{q-1}(1-q/2)^{(2-q)/q}\epsilon^{1/q}n^{1-1/q}}.$$

D. Proof of Proposition 3.3 on page 19

Let us put

$$\widetilde{N}(\theta,\xi) = \inf\left\{\rho \in \mathbb{R}_+^*, \sum_{i=1}^n \psi\left[\lambda\left(\frac{\big(\langle\theta,X\rangle-\xi\big)^2}{\rho}-1\right)\right] \leq 0\right\}.$$

From Lemma 3.2 on page 19, for any $(\theta,\xi) \in \mathbb{R}^{d+1}$,

$$\mathbb{E}\big[\big(\langle\theta,X\rangle-\xi\big)^4\big] \leq (\kappa^{1/2}+1)^2\,\mathbb{E}\big[\big(\langle\theta,X\rangle-\xi\big)^2\big]^2.$$

We can therefore apply Proposition 1.1 on page 7 and obtain for any sample size n satisfying equation (10) on page 19 that with probability at least $1-2\epsilon$, for any $(\theta,\xi)\in\mathbb{R}^{d+1}$,

$$\left|\frac{\mathbb{E}\big[\big(\langle\theta,x\rangle-\xi\big)^2\big]}{\widetilde{N}(\theta,\xi)}-1\right|\leq\frac{\mu}{1-2\mu}.$$

We can then apply Proposition 3.1 on page 18 to conclude.

E. Proof of lemma 3.4 on page 20

Consider $a=\mathbb{E}\big(\langle\theta,X-\mathbb{E}(X)\rangle^2\big)^{1/2}$. If $a=0$, then almost surely $\langle\theta,X_i-\mathbb{E}(X)\rangle=0$, so that in this case

$$\begin{aligned}&\langle\theta,X_i\rangle-\xi=\langle\theta,\mathbb{E}(X)\rangle-\xi,\\ \text{and}\qquad&\mathbb{E}\big[\big(\langle\theta,X\rangle-\xi\big)^2\big]=\big(\langle\theta,\mathbb{E}(X)\rangle-\xi\big)^2\leq 1.\end{aligned}$$

Consequently $\langle\theta,X_i\rangle-\xi\leq 1$.

Assume now that $a>0$. In this case, from the Cauchy-Schwarz inequality in $\mathbb{R}^2$,

$$\begin{aligned}\big(\langle\theta,X_i\rangle-\xi\big)^2&=\big(aa^{-1}\langle\theta,X_i-\mathbb{E}(X)\rangle+\langle\theta,\mathbb{E}(X)\rangle-\xi\big)^2\\&\leq\big[a^{-2}\langle\theta,X_i-\mathbb{E}(X)\rangle^2+1\big]\big[a^2+\big(\langle\theta,\mathbb{E}(X)\rangle-\xi\big)^2\big].\end{aligned}$$

Remark that almost surely $X_i-\mathbb{E}(X)\in\mathbf{Im}(\Sigma)$, so that $\big(X_i-\mathbb{E}(X)\big)=\Sigma^{1/2}\Sigma^{-1/2}\big(X_i-\mathbb{E}(X)\big)$, and therefore,

$$\begin{aligned}\langle\theta,X_i-\mathbb{E}(X)\rangle^2&=\langle\Sigma^{1/2}\theta,\Sigma^{-1/2}\big(X_i-\mathbb{E}(X)\rangle^2\\&\leq\|\Sigma^{1/2}\theta\|^2\|\Sigma^{-1/2}(X_i-\mathbb{E}(X))\|^2=a^2\|\Sigma^{-1/2}\big(X_i-\mathbb{E}(X)\big)\|^2.\end{aligned}$$

On the other hand,

$$\begin{aligned}a^2+\big(\langle\theta,\mathbb{E}(X)-\xi\big)^2&=\mathbb{E}\big(\langle\theta,X-\mathbb{E}(X)\rangle^2\big)+\big(\langle\theta,\mathbb{E}(X)-\xi\big)^2\\&=\mathbb{E}\big[\big(\langle\theta,X\rangle-\xi\big)^2\big]\leq 1.\end{aligned}$$

Putting the three last equations together, we obtain that

$$\big(\langle\theta,X_i\rangle-\xi\big)^2\leq\Big(\big\|\Sigma^{-1/2}\big(X_i-\mathbb{E}(X)\big)\big\|^2+1\Big),$$

which achieves the proof.

F. Proof of Proposition 4.9 on page 34

Let us put in this proof

$$R(\theta) = \mathbb{E}\big[\big(\widetilde{Y} - \langle \theta, \widetilde{X}\rangle\big)^2\big] = p^{-1}\mathbb{E}\big[\big(Y - \langle \theta, X\rangle\big)^2\big].$$

Let $\eta_1, \dots, \eta_n$ be n independent copies of η, independent from $\widetilde{X}_1, \dots, \widetilde{X}_n$ and such that $\widetilde{Y}_i = \langle \theta_*, \widetilde{X}_i\rangle + \eta_i$. Consider

$$\widetilde{G} = \frac{1}{n}\sum_{i=1}^n \widetilde{X}_i \widetilde{X}_i^\top, \qquad W = \frac{1}{\sqrt{n}\sigma}\sum_{i=1}^n \eta_i \widetilde{G}^{-1/2}\widetilde{X}_i, \quad \text{and} \quad \widetilde{d} = \mathbf{rank}(\widetilde{G}).$$

Remark that

$$\mathbb{P}_{W \,|\, \widetilde{X}_1, \dots, \widetilde{X}_n} = \mathcal{N}\big(0, \widetilde{G}^{-1}\widetilde{G}\,\big),$$

where $\widetilde{G}^{-1}$ is the pseudo inverse of $\widetilde{G}$, so that $\widetilde{G}^{-1}\widetilde{G}$ is the orthogonal projection on $\mathbf{Im}(\widetilde{G})$. Applying Chernoff's bound, we get

$$\mathbb{P}\left(\|W\|^2 \geq \frac{2}{\alpha}\left[\frac{\widetilde{d}}{2}\log(1+\alpha) - \log(\epsilon^{-1})\right] \,\middle|\, \widetilde{X}_1, \dots, \widetilde{X}_n\right) \geq 1 - \epsilon,$$

and therefore, integrating with respect to $\widetilde{X}_1, \dots, \widetilde{X}_n$,

$$\mathbb{P}\left(\|W\|^2 \geq \frac{2}{\alpha}\left[\frac{\widetilde{d}}{2}\log(1+\alpha) - \log(\epsilon^{-1})\right]\right) \geq 1 - \epsilon. \tag{29}$$

Remark now that

$$\widetilde{\theta} = \widetilde{G}^{-1}\left(\frac{1}{n}\sum_{i=1}^n \widetilde{Y}_i \widetilde{X}_i\right) + \xi,$$

where $\xi \in \mathbf{Ker}(\widetilde{G})$. Therefore

$$\widetilde{\theta} = \widetilde{G}^{-1}\left(\frac{1}{n}\sum_{i=1}^n \langle \theta_*, X_i\rangle X_i + \eta_i X_i\right) + \xi = \widetilde{G}^{-1}\widetilde{G}\theta_* + \frac{1}{n}\sum_{i=1}^n \eta_i \widetilde{G}^{-1}\widetilde{X}_i + \xi,$$

We see that

$$\begin{aligned}\frac{1}{n}\sum_{i=1}^n \langle \widetilde{\theta} - \theta_*, \widetilde{X}_i\rangle^2 = \big(\widetilde{\theta} - \theta_*\big)^\top \widetilde{G}\big(\widetilde{\theta} - \theta_*\big) &= \big\|\widetilde{G}^{1/2}\big(\widetilde{\theta} - \theta_*\big)\big\|^2 \\ &= \left\|\frac{1}{n}\sum_{i=1}^n \eta_i \widetilde{G}^{-1/2}\widetilde{X}_i\right\|^2 = \frac{\sigma^2}{n}\|W\|^2.\end{aligned}$$

When $n \geq n_\epsilon$, with probability at least $1-2\epsilon$, equation (19) on page 34 is satisfied, so that $\widetilde{d} = \mathbf{rank}(\widetilde{G}) = \mathbf{rank}(G) = d$ and

$$R(\widetilde{\theta}) - R(\theta_*) = \mathbb{E}\big(\langle \widetilde{\theta} - \theta_*, X\rangle^2\big) \geq \frac{1}{(1+\delta)\, n} \sum_{i=1}^{n} \langle \widetilde{\theta} - \theta_*, X_i\rangle^2 = \frac{\sigma^2}{(1+\delta)\, n} \|W\|^2.$$

Combining this inequality with equation (29), we obtain that with probability at least $1-3\epsilon$,

$$R(\widetilde{\theta}) - R(\theta_*) \geq \frac{\sigma^2}{(1+\delta)\, \alpha n} \big[d \log(1+\alpha) - 2\log(\epsilon^{-1}) \big].$$

Choosing for simplicity $\alpha = 1$ gives the first statement of the proposition.

Consider $\sigma_1, \ldots, \sigma_n$, n independent copies of σ, independent of

$$(\widetilde{X}_1, \widetilde{Y}_1), \ldots, (\widetilde{X}_n, \widetilde{Y}_n),$$

and such that $(X_i, Y_i) = \sigma_i(\widetilde{X}_i, \widetilde{Y}_i)$. Define

$$\overline{p} = \frac{1}{n} \sum_{i=1}^{n} \sigma_i.$$

Conditioning with respect to $\sigma_1, \ldots, \sigma_n$, we deduce from the first part of the proposition that

$$\mathbb{P}\left(R(\widehat{\theta}) - R(\theta_*) \geq \frac{\sigma^2 \big[d \log(2) - 2\log(\epsilon^{-1}) \big]}{(1+\delta)\overline{p} n} \,\Big|\, n_\epsilon \leq n\overline{p} \leq 7np/4 \right) \geq 1 - 3\epsilon.$$

We proved in [11, Proposition 20.1 page 289] that with probability at least $1-2\epsilon$,

$$\big|\overline{p} - p\big| \leq \sqrt{2\log(\epsilon^{-1})p/n} + 2\log(\epsilon^{-1})/n.$$

Therefore, when

$$n \geq \frac{8\log(\epsilon^{-1})}{p} = \frac{8(\kappa_1+\kappa_2)\log(\epsilon^{-1})}{\widetilde{\kappa}_1 + \widetilde{\kappa}_2},$$

with probability at least $1-2\epsilon$,

$$p/4 \leq \overline{p} \leq 7p/4.$$

When moreover

$$n \geq \frac{4n_\epsilon}{p} = \frac{4(\kappa_1+\kappa_2)n_\epsilon}{\widetilde{\kappa}_1 + \widetilde{\kappa}_2}.$$

as required in the proposition, with probability at least $1-2\epsilon$,

$$n_\epsilon \leq n\overline{p} \leq 7np/4.$$

We can then write that

$$\begin{aligned}
\mathbb{P}\Big(R(\widehat{\theta}) - R(\theta_*) \geq \frac{4\sigma^2\big[d\log(2) - 2\log(\epsilon^{-1})\big]}{7(1+\delta)pn}\Big) & \\
\geq \mathbb{P}\Big(R(\widehat{\theta}) - R(\theta_*) \geq \frac{4\sigma^2\big[d\log(2) - 2\log(\epsilon^{-1})\big]}{7(1+\delta)pn} & \\
\text{and } n_\epsilon \leq n\overline{p} \leq 7np/4\Big) & \\
= \mathbb{P}\Big(R(\widehat{\theta}) - R(\theta_*) \geq \frac{4\sigma^2\big[d\log(2) - 2\log(\epsilon^{-1})\big]}{7(1+\delta)pn} \,\Big|\, n_\epsilon \leq n\overline{p} \leq 7np/4\Big) & \\
\times \mathbb{P}\Big(n_\epsilon \leq n\overline{p} \leq 7np/4\Big) & \\
\geq \mathbb{P}\Big(R(\widehat{\theta}) - R(\theta_*) \geq \frac{\sigma^2\big[d\log(2) - 2\log(\epsilon^{-1})\big]}{(1+\delta)\overline{p}n} \,\Big|\, n_\epsilon \leq n\overline{p} \leq 7np/4\Big)(1-2\epsilon) & \\
\geq (1-3\epsilon)(1-2\epsilon) \geq 1-5\epsilon. &
\end{aligned}$$

We conclude the proof of equation (20) on page 35 by remarking that

$$\mathbb{E}\big[\big(Y - \langle\theta, X\rangle\big)^2\big] = pR(\theta) \text{ and } \mathbb{E}\big[\big(Y - \langle\theta_*, X\rangle\big)^2\big] = pR(\theta_*) = p\sigma^2,$$

and that

$$p^{-1} = \frac{\kappa_1 + \kappa_2}{\widetilde{\kappa}_1 + \widetilde{\kappa}_2}.$$

Applying Proposition 2.3 on page 14 to a Gaussian design, we obtain, for any parameter $\alpha \in]0,1[$, that

$$\widehat{\gamma}_+ \leq \frac{(1+\widehat{\delta})^2}{3} a\Big[\frac{b}{\sqrt{n}} + \frac{2\log(n)}{\alpha\sqrt{n}}\Big]^2,$$

where

$$\begin{aligned}
a &= 0.73\, d + \log(\epsilon^{-1}), \\
b &= \frac{d}{\alpha}\log\Big(\frac{1}{1-\alpha}\Big) + \frac{2}{\alpha}\log(\epsilon^{-1}).
\end{aligned}$$

We can then use the fact that the log is concave and that for any constant $c > 0$

$$\log(n) = 3\log\big(n^{1/3}\big) \leq 3\Big(\log(c) + \frac{n^{1/3} - c}{c}\Big) = 3\Big(\log(c/e) + \frac{n^{1/3}}{c}\Big).$$

Therefore, $\widehat{\gamma}_+ \leq \chi$ when

$$\left(\frac{b + \frac{6}{\alpha} \log\bigl(c/e\bigr)}{n^{1/2}} + \frac{6}{\alpha c n^{1/6}} \right)^2 \leq \frac{3\chi}{(1+\widehat{\delta})^2 a}.$$

This condition is satisfied when

$$\frac{b + \frac{6}{\alpha} \log\bigl(c/e\bigr)}{n^{1/2}} \leq \frac{1}{2} \sqrt{\frac{3\chi}{(1+\widehat{\delta})^2 a}},$$
$$\text{and } \frac{6}{\alpha c n^{1/6}} \leq \frac{1}{2} \sqrt{\frac{3\chi}{(1+\widehat{\delta})^2 a}}.$$

This can be rewritten as

$$n^{1/2} \geq 2 \frac{(1+\widehat{\delta})}{(3\chi)^{1/2}} \Bigl(b + 6 \log\bigl(c/e\bigr)/\alpha \Bigr) a^{1/2},$$
$$\text{and } n^{1/2} \geq \left(\frac{12}{\alpha c} \right)^3 \frac{(1+\widehat{\delta})^3}{\bigl(3\chi\bigr)^{3/2}} a^{3/2}.$$

Due to the fact that $b \geq a$, the first condition implies the second one when

$$c^3 = \frac{1}{2} \left(\frac{12}{\alpha} \right)^3 \frac{\bigl(1+\widehat{\delta}\bigr)^2}{3\chi},$$

so that with this choice of constant c the condition becomes

$$n_\epsilon \geq 4 \frac{(1+\widehat{\delta})^2}{3\chi} a \bigl[b + 6 \log\bigl(c/e\bigr)/\alpha \bigr]^2.$$

Taking $\chi = 1/4$, $\alpha = 1/2$, $\widehat{\delta} = 1/4$ and working out the constants numerically gives the condition on n_ϵ stated in the proposition. One can then check that the condition on n_ϵ necessary to obtain $\mu \leq 1/6$ and therefore $\widehat{\delta} \leq 1/4$ in Proposition 2.3 on page 14 is weaker, so that we are entitled to take $\widehat{\delta} \leq 1/4$, which gives $\widehat{\gamma}_+ \leq \chi = 1/4$ and

$$\delta = \frac{\widehat{\delta} + \widehat{\gamma}_+}{(1-\widehat{\delta})(1-\widehat{\gamma}_+)} \leq \frac{8}{9} < 1$$

as required. The last equation of the proposition is then obtained by substituting this value in equation (20) on page 35.

G. Proof of Proposition 4.10 on page 36

Consider the empirical Gram matrix

$$\overline{G} = \frac{1}{n}\sum_{i=1}^{n} X_i X_i^\top.$$

According to Proposition 2.2 on page 11, under the same assumptions as in Proposition 1.1 on page 7, there is an event Ω of probability at least $1-\epsilon$ on which for any $\theta \in \mathbb{R}^d$,

$$\frac{\theta^\top \overline{G}\theta}{\theta^\top G\theta} \geq 1 - \frac{\widehat{\delta}+\gamma_-}{1+\widehat{\delta}} = \frac{1-\gamma_-}{1+\widehat{\delta}}.$$

(Given that we use only one side of the inequalities involved in Proposition 2.2 and that each side holds with probability at least $1-\epsilon$, as is clear from the proof of this proposition.) Remark that on Ω, $\mathbf{Im}(G) \subset \mathbf{Im}(\overline{G})$. As $\mathbf{Im}(\overline{G}) \subset \mathbf{Im}(G)$ almost surely, we may remove a set of measure zero from Ω and assume that on Ω, $\mathbf{Im}(G) = \mathbf{Im}(\overline{G})$. Since it does not change $R(\theta_*)$, we may assume without loss of generality that $\theta_* \in \mathbf{Im}(G)$, by projecting it on $\mathbf{Im}(G)$ if necessary. Remark then that for some $\xi \in \mathbf{Ker}(G) = \mathbf{Ker}(\overline{G})$,

$$\begin{aligned}\widehat{\theta} = \overline{G}^{-1}\left(\frac{1}{n}\sum_{i=1}^{n} Y_i X_i\right) + \xi &= \overline{G}^{-1}\left(\frac{1}{n}\sum_{i=1}^{n}\big(Y_i - \langle\theta_*, X_i\rangle\big)X_i\right)\\ &\quad + \overline{G}^{-1}\left(\frac{1}{n}\sum_{i=1}^{n}\langle\theta_*, X_i\rangle X_i\right) + \xi = \overline{G}^{-1}W + \theta_* + \xi,\end{aligned}$$

where we have introduced the notation

$$W = \frac{1}{n}\sum_{i=1}^{n}\big(Y_i - \langle\theta_*, X_i\rangle\big)X_i. \tag{30}$$

Let us put $\rho = \dfrac{1+\widehat{\delta}}{1-\gamma_-}$ and let us remark that on the event Ω, for any $\theta \in \mathbf{Im}(G)$,

$$\begin{aligned}\big\|\theta\big\|^2 = \big\|G^{1/2}G^{-1/2}\theta\big\|^2 &\leq \rho\,\big\|\overline{G}^{1/2}G^{-1/2}\theta\big\|^2\\ &= \rho\,\big\|\big(G^{-1/2}\overline{G}G^{-1/2}\big)^{1/2}\theta\big\|^2 \leq \rho^2\big\|G^{-1/2}\,\overline{G}\,G^{-1/2}\,\theta\big\|^2.\end{aligned}$$

Assuming that we are on the event Ω, we can then write

$$R(\widehat{\theta}) - R(\theta_*) = \big\|G^{1/2}\big(\widehat{\theta}-\theta_*\big)\big\|^2 = \big\|G^{1/2}\overline{G}^{-1}W\big\|^2 \leq \rho^2\big\|G^{-1/2}W\big\|^2.$$

Thus,

$$\mathbb{E}\Big[\big(R(\widehat{\theta})-R(\theta_*)\big)\mathbb{1}_\Omega\Big] \le \rho^2\mathbb{E}\Big[\big\|G^{-1/2}W\big\|^2\mathbb{1}_\Omega\Big] \le \rho^2\mathbb{E}\Big[\big\|G^{-1/2}W\big\|^2\Big] \\ = \frac{(1+\widehat{\delta})^2}{(1-\gamma_-)^2\,n}\mathbb{E}\Big[\big(Y-\langle\theta_*,X\rangle\big)^2\big\|G^{-1/2}X\big\|^2\Big].$$

This proves the second statement of the proposition. The first statement is deduced from the second one, writing that

$$\mathbb{E}\Big[\min\big\{R(\widehat{\theta}-R(\theta_*),M\big\}\Big] \le \mathbb{E}\Big[M\mathbb{1}_{\Omega^c}+\big[R(\widehat{\theta})-R(\theta_*)\big]\mathbb{1}_\Omega\Big] \\ \le M\epsilon+\mathbb{E}\Big[\big(R(\widehat{\theta})-R(\theta_*)\big)\mathbb{1}_\Omega\Big].$$

Then if we take $\epsilon = C/(Mn^2)$, the condition $n \ge \mathbf{O}\Big(\kappa\big[d+\log(\epsilon^{-1})\big]\Big)$ becomes $n \ge a\kappa\big[d+\log\big(Mn^2/C\big)\big]$, for some constant $a>0$. Consider the constant $b = 4a\kappa$. Remarking that $\log(n) \le \log(b/e)+n/b$, we get that

$$a\kappa\big[d+\log\big(Mn^2/C\big)\big] \le a\kappa\big[d+\log\big(M/C\big)+2\log\big(b/e\big)+2n/b\big] \\ = a\kappa\big[d+\log\big(M/C\big)+2\log\big(4a\kappa/e\big)\big]+n/2,$$

so that the condition above is implied by the condition

$$n \ge 2a\kappa\big[d+\log\big(M/C\big)+2\log\big(4a\kappa/e\big)\big] = \mathbf{O}\Big(\kappa\big[d+\log\big(\kappa M/C\big)\big]\Big). \quad (31)$$

Let us now assume that $Y = \langle\theta_*,X\rangle+\eta$, where η is independent from X and such that $\mathbb{E}(\eta)=0$ and $\mathbb{E}\big(\sigma^2\big)=\sigma^2$. The event Ω described above is measurable with respect to the sigma-algebra generated by $(X_1,\dots,X_n)$. Reasoning as previously we see that on Ω,

$$\mathbb{E}\big[R(\widehat{\theta})-R(\theta_*)\;\big|\;X_1,\dots,X_n\big] \le \rho\,\mathbb{E}\Big[\big\|\overline{G}^{-1/2}W\big\|^2\;\big|\;X_1,\dots,X_n\Big] \\ = \rho\frac{\mathbf{rank}(\overline{G})\sigma^2}{n} \le \frac{\rho d\sigma^2}{n}.$$

The statement about $\mathbb{E}\Big[\min\big\{R(\widehat{\theta})-R(\theta_*),M\big\}\Big]$ is then deduced as above.

H. Proof of Proposition 4.12 on page 38

From Proposition 2.2 on page 11, we see that, when $n \ge n_\epsilon$ given by equation 3 on page 7 on some event Ω of probability at least $1-2\epsilon$, for any

$\theta \in \mathbb{R}^d$,

$$(1-\widehat{\delta})(1-\widetilde{\gamma}_+)\|\overline{G}^{1/2}\theta\|^2 \leq \|G^{1/2}\theta\|^2 \leq \frac{1+\widehat{\delta}}{1-\gamma_-}\|\overline{G}^{1/2}\theta\|^2.$$

Let us put for short $\rho_+ = \dfrac{1+\widehat{\delta}}{1-\gamma_-}$ and $\rho_- = (1-\widehat{\delta})(1-\widetilde{\gamma}_+)$.

As already shown in a previous proof, the upper bound implies that $\mathbf{Im}(\overline{G}) = \mathbf{Im}(G)$. We can for this reason write that, on Ω,

$$R(\widehat{\theta}) - R(\theta_*) = \big\|G^{1/2}\big(\widehat{\theta}-\theta_*\big)\big\|^2 = \big\|G^{1/2}\overline{G}^{-1}W\big\|^2,$$

where the random variable W is defined by equation (30) on page 63. Moreover, for any $\theta \in \mathbf{Im}(G)$,

$$\|\theta\|^2 = \|G^{1/2}G^{-1/2}\theta\|^2 \geq \rho_-\,\|\overline{G}^{1/2}G^{-1/2}\theta\|^2 = \rho_-\,\big\|\big(G^{-1/2}\overline{G}G^{-1/2}\big)^{1/2}\theta\big\|^2.$$

Iterate this inequality to get, for any $\theta \in \mathbf{Im}(G)$,

$$\|\theta\|^2 \geq \rho_-^2\big\|G^{-1/2}\overline{G}G^{1/2}\theta\|^2.$$

Apply this to $\theta = G^{1/2}\overline{G}^{-1}W$, to get on the event Ω

$$R(\widehat{\theta}) - R(\theta_*) \geq \rho_-^2\big\|G^{-1/2}W\big\|^2.$$

We have also a reverse inequality on Ω, that is proved as in Propsition 4.10 on page 36 (where the event Ω was larger), and writes as

$$R(\widehat{\theta}) - R(\theta_*) \leq \rho_+^2\big\|G^{-1/2}W\big\|^2.$$

Consequently $\mathbb{E}\Big(\big[R(\widehat{\theta}) - R(\theta_*)\big]\mathbb{1}_\Omega\Big) \leq \rho_+^2\mathbb{E}\Big(\big\|G^{-1/2}W\big\|^2\mathbb{1}_\Omega\Big)$

$$\leq \rho_+^2\mathbb{E}\Big(\big\|G^{-1/2}W\big\|^2\Big)$$

and $\mathbb{E}\Big(\big[R(\widehat{\theta}) - R(\theta_*)\big]\mathbb{1}_\Omega\Big) \geq \rho_-^2\mathbb{E}\Big(\big\|G^{-1/2}W\big\|^2\mathbb{1}_\Omega\Big)$

$$= \rho_-^2\Big[\mathbb{E}\Big(\big\|G^{-1/2}W\big\|^2\Big) - \mathbb{E}\Big(\big\|G^{-1/2}W\big\|^2\mathbb{1}_{\Omega^c}\Big)\Big]$$

$$\geq \rho_-^2\Big[\mathbb{E}\Big(\big\|G^{-1/2}W\big\|^2\Big) - \mathbb{E}\Big(\big\|G^{-1/2}W\big\|^4\Big)^{1/2}\mathbb{P}\big(\Omega^c\big)^{1/2}\Big],$$

where we have used the Cauchy-Schwarz inequality. Let us put

$$G^{-1/2}W = \frac{1}{n}\sum_{i=1}^{n} Z_i,$$

where

$$Z_i = \big(Y_i - \langle \theta_*, X_i\rangle\big) G^{-1/2} X_i.$$

Remark that $\mathbb{E}(Z_i) = 0$, $\mathbb{E}\big(\|Z_i\|^2\big) = C$, $\mathbb{E}\big(\|Z_i\|^4\big) = \kappa' C^2$ and that the random vectors Z_i are i.i.d. Compute

$$\begin{aligned}
\mathbb{E}\Big(\big\|G^{-1/2}W\big\|^2\Big) &= \mathbb{E}\Big(\big\langle G^{-1/2}W, G^{-1/2}W\big\rangle\Big) = \frac{1}{n^2}\sum_{i=1}^{n}\mathbb{E}\Big(\|Z_i\|^2\Big) = \frac{C}{n},\\
\mathbb{E}\Big(\big\|G^{-1/2}W\big\|^4\Big) &= \mathbb{E}\Big(\big\langle G^{-1/2}W, G^{-1/2}W\big\rangle^2\Big)\\
&= \frac{1}{n^4}\mathbb{E}\left[\left(\sum_{i=1}^{n}\|Z_i\|^2 + 2\sum_{1\le i<j\le n}\langle Z_i, Z_j\rangle\right)^2\right]\\
&= \frac{1}{n^4}\mathbb{E}\left(\sum_{i=1}^{n}\|Z_i\|^4 + 2\sum_{1\le i<j\le n}\|Z_i\|^2\|Z_j\|^2 + 4\sum_{1\le i<j\le n}\langle Z_i, Z_j\rangle^2\right)\\
&\le \frac{1}{n^4}\mathbb{E}\left(\sum_{i=1}^{n}\|Z_i\|^4 + 2\sum_{1\le i<j\le n}\|Z_i\|^2\|Z_j\|^2 + 4\sum_{1\le i<j\le n}\|Z_i\|^2\|Z_j\|^2\right)\\
&= \frac{n\kappa' + 3n(n-1)}{n^4}\,C^2 = 3\left(1 + \frac{\kappa'-3}{3n}\right)\frac{C^2}{n^2}.
\end{aligned}$$

This proves that

$$\begin{aligned}
\rho_-^2\left[1 - \sqrt{6}\left(1 + \frac{\kappa'-3}{3n}\right)^{1/2}\epsilon^{1/2}\right]\frac{C}{n} &\le \mathbb{E}\Big(\big[R(\widehat{\theta}) - R(\theta_*)\big]\mathbb{1}_\Omega\Big)\\
&\le \mathbb{E}\Big(R(\widehat{\theta}) - R(\theta_*)\,|\,\Omega\Big) \le \frac{\rho_+^2\, C}{n\,\mathbb{P}(\Omega)},
\end{aligned}$$

as stated in equation (21) on page 39 of the proposition. Now

$$\begin{aligned}
\mathbb{E}\Big(\min\big\{R(\widehat{\theta}) - R(\theta_*), M\big\}\Big) &\ge \mathbb{E}\Big(\min\big\{R(\widehat{\theta}) - R(\theta_*), M\big\}\mathbb{1}_\Omega\Big)\\
&\ge \mathbb{E}\Big(\min\big\{\rho_-^2\big\|G^{-1/2}W\big\|^2, M\big\}\mathbb{1}_\Omega\Big)\\
&\ge \rho_-^2\,\mathbb{E}\Big(\min\big\{\big\|G^{-1/2}W\big\|^2, M\big\}\mathbb{1}_\Omega\Big)
\end{aligned}$$

$$\geq \rho_-^2 \left[\mathbb{E}\Big(\min\big\{ \big\| G^{-1/2} W \big\|^2 , M \big\} \Big) - \mathbb{E}\Big(\big\| G^{-1/2} W \big\|^2 \mathbb{1}_{\Omega^c} \Big) \right]$$

$$\geq \rho_-^2 \left[\mathbb{E}\Big(\big\| G^{-1/2} W \big\|^2 \Big) - \mathbb{E}\Big[\Big(\big\| G^{-1/2} W \big\|^2 - M \Big)_+ \Big] \right.$$

$$\left. - \mathbb{E}\Big(\big\| G^{-1/2} W \big\|^4 \Big)^{1/2} \mathbb{P}\big(\Omega^c\big)^{1/2} \right].$$

Remark that for any $z \in \mathbb{R}$, $(z - M)_+ \leq z^2/(4M)$. Thus

$$\mathbb{E}\Big[\Big(\big\| G^{-1/2} W \big\|^2 - M \Big)_+ \Big] \leq \frac{1}{4M} \mathbb{E}\Big(\big\| G^{-1/2} W \big\|^4 \Big) = \frac{3}{4M} \left(1 + \frac{\kappa' - 3}{3n} \right) \frac{C^2}{n^2}$$

and

$$\mathbb{E}\Big(\min\big\{ R(\widehat{\theta}) - R(\theta_*), M \big\} \Big)$$

$$\geq \rho_-^2 \left[\frac{C}{n} - \frac{3}{4M} \left(1 + \frac{\kappa' - 3}{3n} \right) \frac{C^2}{n^2} - \frac{\sqrt{6} C}{n} \left(1 + \frac{\kappa' - 3}{3n} \right)^{1/2} \epsilon^{1/2} \right],$$

that proves equation (22) of the proposition. The end of the proposition is straightforward, the evaluation of n_ϵ in big **O** notation when $\epsilon = n^{-(q-1)}$ being done using the same principle as in the proof of equation (31) on page 64.

I. Obtaining a quadratic form

In this section, we will see how to deduce from the estimator of Proposition 1.1 on page 7, that is not a quadratic form, a quadratic estimator, or equivalently an estimator of the Gram matrix $G = \mathbb{E}\big(XX^\top\big)$ by a symmetric non-negative matrix $\widehat{G}$.

Let us assume that we derived as in Proposition 1.1 on page 7 an estimator $\widehat{N}$ such that for some ϵ and $\delta \in]0, 1/2[$, on some event Ω' of probability at least $1 - \epsilon$, for any $\theta \in \mathbb{R}^d$,

$$\left| \frac{N(\theta)}{\widehat{N}(\theta)} - 1 \right| \leq \delta. \tag{32}$$

Assume moreover that, as it is the case in Proposition 1.1, for any $\theta \in \mathbb{R}^d$ such that $\theta \perp \mathbf{span}\big\{ X_1, \ldots, X_n \big\}$, $\widehat{N}(\theta) = 0$.

Let us remark that equation (32) implies that on Ω'

$$\mathbf{Ker}(G) = \big\{\theta \in \mathbb{R}^d \,:\, N(\theta) = 0\big\} = \big\{\theta \in \mathbb{R}^d \,:\, \widehat{N}(\theta) = 0\big\}.$$

Remark then that our second assumption implies that on Ω'

$$\mathbf{span}\big\{X_1, \dots, X_n\big\}^\perp \subset \mathbf{Ker}(G).$$

Remark moreover that almost surely

$$\mathbf{Ker}(G) \subset \mathbf{span}\big\{X_1, \dots, X_n\big\}^\perp,$$

since for any θ in a basis of $\mathbf{Ker}(G)$ (that is a finite set), $\mathbb{E}\big(\langle \theta, X\rangle^2\big) = 0$, so that almost surely $\langle X_i, \theta\rangle^2 = 0$, $1 \le i \le n$.

This proves that under our assumptions, on some event Ω'' of the same probability as Ω',

$$\mathbf{Im}(G) = \mathbf{span}\big\{X_1, \dots, X_n\big\},$$

so that we have an easily computable estimator of $\mathbf{Im}(G)$ that is exact on Ω'', an event of probability at least $1 - \epsilon$.

Consider some positive parameter ρ, let $\mathbb{S}_d = \big\{\theta \in \mathbb{R}^d \,:\, \|\theta\| = 1\big\}$ be the unit sphere of $\mathbb{R}^d$, and Θ_ρ some arbitrary ρ-net of $\mathbf{span}\{X_1, \dots, X_n\} \cap \mathbb{S}_d$. In other words, assume that Θ_ρ is a finite subset of $\mathbf{span}\{X_1, \dots, X_n\} \cap \mathbb{S}_d$ such that

$$\sup\Big\{ \inf_{\xi \in \Theta_\rho} \|\theta - \xi\| \,:\, \theta \in \mathbf{span}\big\{X_1, \dots, X_n\big\} \cap \mathbb{S}_d \Big\} \le \rho.$$

Assume now that $\widehat{G} \in \mathbb{R}^{d\times d}$ is a random symmetric matrix solution of

$$\widehat{G} = \arg\min_H \Big\{ \mathbf{Tr}(H^2) \,:\, H \in \mathbb{R}^{d\times d}, H^\top = H, \\ \mathbf{Im}(H) \subset \mathbf{span}\big\{X_1, \dots, X_n\big\}, \\ \widehat{N}(\theta)(1-\delta) \le \theta^\top H\, \theta \le \widehat{N}(\theta)(1+\delta), \quad \theta \in \Theta_\rho \Big\}.$$

This minimization problem has a solution on Ω'', since in this case G itself satisfies the constraints. We will see below in more detail that $\widehat{G}$ is the solution of a convex minimization problem very similar to the one appearing in the estimation of the parameters of a support vector machine using the popular box constraint learning algorithm.

Proposition I.1 *The symmetric matrix $\widehat{G}$ is such that on the event Ω'' of probability at least $1 - \epsilon$, $\mathbf{Im}(\widehat{G}) \subset \mathbf{Im}(G)$ and for any $\theta \in \mathbb{R}^d$,*

$$\Big|\theta^\top \widehat{G}\theta - N(\theta)\Big| \le \frac{2\delta}{1-\delta} N(\theta) + \frac{4\rho\sqrt{\mathbf{Tr}(G^2)}}{(1-\delta)} \|\theta\|^2.$$

The positive part $\widehat{G}_+$ of $\widehat{G}$ is such that on the event Ω'' of probability at least $1-\epsilon$, for any $\theta \in \mathbb{R}^d$,

$$\left|\theta^\top \widehat{G}_+ \theta - N(\theta)\right| \leq \frac{2\delta}{1-\delta} N(\theta) + \frac{6\rho\sqrt{\mathbf{T}\mathrm{r}(G^2)}}{(1-\delta)} \|\theta\|^2.$$

Proof. During all this proof, we will assume that the event Ω'' defined above is satisfied, so that the results will hold with probability at least $1-\epsilon$.

Let us also assume first that $\theta \in \mathbf{Im}(G) \cap \mathbb{S}_d$. Recall that on Ω'', $\mathbf{Im}(G) = \mathbf{span}\{X_1, \ldots, X_n\}$, so that by construction of $\widehat{G}$, $\mathbf{Im}(\widehat{G}) \subset \mathbf{Im}(G)$.

Since Θ_ρ is a ρ-net of $\mathbf{Im}(G) \cap \mathbb{S}_d$, there is ξ in Θ_ρ such that $\|\theta - \xi\| \leq \rho$. Consequently,

$$\left|\theta^\top \widehat{G} \theta - \xi^\top \widehat{G} \xi\right| = \left|(\theta+\xi)^\top \widehat{G} (\theta - \xi)\right| \leq 2\rho \, \|\widehat{G}\|_\infty,$$

where

$$\|\widehat{G}\|_\infty = \sup\left\{ \theta \in \mathbb{S}_d \, : \, \|\widehat{G}\theta\| \right\}$$

is the operator norm—and spectral radius, since $\widehat{G}$ is symmetric—of $\widehat{G}$. Moreover,

$$\begin{aligned}
\left|\xi^\top \widehat{G} \xi - N(\theta)\right| &\leq \left|\xi^\top \widehat{G} \xi - \widehat{N}(\xi)\right| + \left|\widehat{N}(\xi) - N(\xi)\right| + \left|\xi^\top G \xi - \theta^\top G \theta\right| \\
&\leq 2\delta \widehat{N}(\xi) + 2\rho \|G\|_\infty \leq \frac{2\delta}{1-\delta} \xi^\top G \xi + 2\rho \|G\|_\infty \\
&\leq \frac{2\delta}{1-\delta}\Big(\theta^\top G \theta + 2\rho \|G\|_\infty\Big) + 2\rho \|G\|_\infty \\
&= \frac{2\delta}{1-\delta} N(\theta) + \frac{2(1+\delta)}{1-\delta} \rho \|G\|_\infty.
\end{aligned}$$

Remark now that $\|\widehat{G}\|_\infty \leq \sqrt{\mathbf{T}\mathrm{r}(\widehat{G}^2)} \leq \sqrt{\mathbf{T}\mathrm{r}(G^2)}$, so that we can deduce from the two above inequalities that

$$\left|\theta^\top \widehat{G} \theta - N(\theta)\right| \leq \frac{2\delta}{1-\delta} N(\theta) + \frac{4\rho}{1-\delta} \sqrt{\mathbf{T}\mathrm{r}(G^2)}, \qquad \theta \in \mathbf{Im}(G) \cap \mathbb{S}_d.$$

Recall that $\mathbf{Im}(\widehat{G}_-) \subset \mathbf{Ker}(\widehat{G}_+) \cap \mathbf{Im}(G)$, so that for any $\theta \in \mathbf{Im}(\widehat{G}_-) \cap \mathbb{S}_d$, there is $\xi \in \Theta_\rho$ such that $\|\theta - \xi\| \leq \rho$. Consequently

$$\begin{aligned}
\theta^\top \widehat{G}_- \theta = -\theta^\top \widehat{G} \theta &\leq -\xi^\top \widehat{G} \xi + 2\rho \|\widehat{G}\|_\infty \\
&\leq -(1-\delta)\widehat{N}(\xi) + 2\rho\sqrt{\mathbf{T}\mathrm{r}(\widehat{G}^2)} \leq 2\rho\sqrt{\mathbf{T}\mathrm{r}(G^2)}.
\end{aligned}$$

This proves that $\|\widehat{G}_-\|_\infty \le 2\rho\sqrt{\mathbf{Tr}(G^2)}$. As a consequence,

$$\begin{aligned}\left|\theta^\top \widehat{G}_+\theta - N(\theta)\right| &\le \left|\theta^\top \widehat{G}_+\theta - \theta^\top \widehat{G}\theta\right| + \left|\theta^\top \widehat{G}\theta - N(\theta)\right| \\ &= \theta^\top \widehat{G}_-\theta + \left|\theta^\top \widehat{G}\theta - N(\theta)\right| \le \frac{2\delta}{1-\delta}N(\theta) + \frac{6\rho}{1-\delta}\sqrt{\mathbf{Tr}(G^2)}, \\ &\qquad\qquad \theta \in \mathbf{Im}(G)\cap \mathbb{S}_d.\end{aligned}$$

By homogeneity, we get on the event Ω'' that

$$\left|\theta^\top \widehat{G}\theta - N(\theta)\right| \le \frac{2\delta}{1-\delta}N(\theta) + \frac{4\rho}{1-\delta}\sqrt{\mathbf{Tr}(G^2)}\|\theta\|^2, \qquad \theta \in \mathbf{Im}(G),$$
$$\left|\theta^\top \widehat{G}_+\theta - N(\theta)\right| \le \frac{2\delta}{1-\delta}N(\theta) + \frac{6\rho}{1-\delta}\sqrt{\mathbf{Tr}(G^2)}\|\theta\|^2, \qquad \theta \in \mathbf{Im}(G).$$

Let us now deal with the general case of an arbitrary $\theta \in \mathbb{R}^d$. We can decompose it into $\theta = \theta_1 + \theta_2$, where $\theta_1 \in \mathbf{Im}(G)$ and $\theta_2 \in \mathbf{Ker}(G)$. Since θ_1 and θ_2 are orthogonal, and since $\mathbf{Im}(\widehat{G}_+) \subset \mathbf{Im}(\widehat{G}) \subset \mathbf{Im}(G)$ on Ω'', $\theta_2 \in \mathbf{Ker}(\widehat{G}) \subset \mathbf{Ker}(\widehat{G}_+)$, so that

$$\theta^\top \widehat{G}\theta = \theta_1 \widehat{G}\theta_1, \quad \theta^\top \widehat{G}_+\theta = \theta_1^\top \widehat{G}_+\theta_1, \text{ and } \theta^\top G\theta = \theta_1^\top G\theta_1.$$

Therefore, on the event Ω'' of probability at least $1-\epsilon$,

$$\begin{aligned}\left|\theta^\top \widehat{G}_+\theta - N(\theta)\right| &= \left|\theta_1^\top \widehat{G}_+\theta_1 - N(\theta_1)\right| \\ &\le \frac{2\delta}{1-\delta}N(\theta_1) + \frac{6\rho}{1-\delta}\sqrt{\mathbf{Tr}(G^2)}\|\theta_1\|^2 \\ &= \frac{2\delta}{1-\delta}N(\theta) + \frac{6\rho}{1-\delta}\sqrt{\mathbf{Tr}(G^2)}\big(\|\theta\|^2 - \|\theta_2\|^2\big) \\ &\le \frac{2\delta}{1-\delta}N(\theta) + \frac{6\rho}{1-\delta}\sqrt{\mathbf{Tr}(G^2)}\|\theta\|^2, \qquad \theta \in \mathbb{R}^d,\end{aligned}$$

and, due to a similar chain of inequalities,

$$\left|\theta^\top \widehat{G}\theta - N(\theta)\right| \le \frac{2\delta}{1-\delta}N(\theta) + \frac{4\rho}{(1-\delta)}\sqrt{\mathbf{Tr}(G^2)}\|\theta\|^2, \qquad \theta \in \mathbb{R}^d.$$

□

COROLLARY I.2 *Introduce* $\lambda_{\min} = \inf\Big\{N(\theta) \,:\, \theta \in \mathbf{Im}(G)\cap \mathbb{S}_d\Big\} > 0$, *the smallest non zero eigenvalue of* G. *On the event* Ω'' *of probability at least* $1-\epsilon$, *for any* $\theta \in \mathbb{R}^d$,

$$\left|\frac{\theta^\top \widehat{G}\theta}{\theta^\top G\theta} - 1\right| \le \frac{1}{1-\delta}\left(2\delta + \frac{4\rho\sqrt{\mathbf{Tr}(G^2)}}{\lambda_{\min}}\right).$$

As a consequence, when

$$\frac{1}{1-\delta}\left(2\delta + \frac{4\rho\sqrt{\mathbf{T}\mathrm{r}(G^2)}}{\lambda_{\min}}\right) < 1,$$

that is when

$$\rho < \frac{(1-3\delta)\lambda_{\min}}{4\sqrt{\mathbf{T}\mathrm{r}(G^2)}},$$

$\widehat{G}_+ = \widehat{G}$ *on* Ω''. *Moreover, if we choose* ρ *small enough, and more precisely such that*

$$\rho \leq \frac{\delta^2 \widehat{\lambda}_{\min}}{2\sum_{i=1}^d \widehat{N}(e_i)},$$

where $(e_1, \ldots, e_d)$ *is some arbitrary orthonormal basis of* $\mathbb{R}^d$, *and*

$$\widehat{\lambda}_{\min} = \inf\Big\{\widehat{N}(\theta) \,:\, \theta \in \mathbf{span}\{X_1, \ldots, X_n\} \cap \mathbb{S}_d\Big\},$$

then on Ω'', *and therefore with probability at least* $1-\epsilon$,

$$\left|\frac{\theta^\top \widehat{G}\theta}{\theta^\top G\theta} - 1\right| \leq \frac{2\delta}{1-2\delta}.$$

Proof. To prove the first inequality of the corollary, consider any $\theta \in \mathbb{R}^d$ and write it as $\theta = \theta_1 + \theta_2$, where $\theta_1 \in \mathbf{Im}(G)$ and $\theta_2 \in \mathbf{Ker}(G)$. If $\theta_1 = 0$, since we saw that on Ω'', $\mathbf{Im}(\widehat{G}) \subset \mathbf{Im}(G)$, then $\theta^\top \widehat{G}\theta = 0$, so that with the convention used throughout this paper

$$\frac{\theta^\top \widehat{G}\theta}{\theta^\top G\theta} = \frac{0}{0} = 1.$$

Otherwise, $\theta^\top G\theta = \theta_1^\top G\theta_1 > 0$, and, as seen in the previous proof, on the event Ω'',

$$\frac{\theta^\top \widehat{G}\theta}{\theta^\top G\theta} = \frac{\theta_1^\top \widehat{G}\theta_1}{\theta_1^\top G\theta_1}.$$

Moreover

$$\|\theta_1\|^2 \leq \lambda_{\min}^{-1} N(\theta_1),$$

so that on Ω'', according to the previous proposition,

$$\left|\frac{\theta^\top \widehat{G}\theta}{\theta^\top G\theta} - 1\right| \leq \frac{2\delta}{1-\delta} + \frac{4\rho\sqrt{\mathbf{T}\mathrm{r}(G^2)}}{1-\delta} \times \frac{\|\theta_1\|^2}{N(\theta_1)} \leq \frac{1}{1-\delta}\left(2\delta + \frac{4\rho\sqrt{\mathbf{T}\mathrm{r}(G^2)}}{\lambda_{\min}}\right).$$

To prove the end of the corollary, remark that on Ω''

$$\widehat{\lambda}_{\min}(1-\delta) \leq \lambda_{\min} \leq \widehat{\lambda}_{\min}(1+\delta)$$

since $\mathbf{span}\{X_1, \ldots, X_n\} = \mathbf{Im}(G)$ on Ω'' and since

$$\lambda_{\min} = \inf\Big\{ N(\theta) \,:\, \theta \in \mathbf{Im}(G) \cap \mathbb{S}_d \Big\}.$$

Remark also that on the event Ω''. for any orthonormal basis $(e_1, \ldots, e_d)$ of $\mathbb{R}^d$,

$$\begin{aligned}\sqrt{\mathbf{T}\mathrm{r}(G^2)} \leq \sqrt{\|G\|_\infty \, \mathbf{T}\mathrm{r}(G)} \leq \mathbf{T}\mathrm{r}(G) &= \sum_{i=1}^d N(e_i) \\ &\leq (1+\delta) \sum_{i=1}^d \widehat{N}(e_i) \leq \frac{(1+\delta)}{(1-\delta)} \, \mathbf{T}\mathrm{r}(G),\end{aligned}$$

where $\|G\|_\infty = \sup_{\theta \in \mathbb{S}_d} \|G\theta\|$ is the operator norm (and spectral radius) of G.

Therefore on Ω'', for any $\theta \in \mathbb{R}^d$,

$$\left| \frac{\theta^\top \widehat{G} \theta}{\theta^\top G \theta} - 1 \right| \leq \frac{1}{1-\delta} \left(2\delta + \frac{4\rho \sum_{i=1}^d \widehat{N}(e_i)(1+\delta)}{\widehat{\lambda}_{\min}(1-\delta)} \right).$$

When $\rho \leq \dfrac{\delta^2 \widehat{\lambda}_{\min}}{2\sum_{i=1}^d \widehat{N}(e_i)}$, we obtain

$$\left| \frac{\theta^\top \widehat{G} \theta}{\theta^\top G \theta} - 1 \right| \leq \frac{1}{1-\delta} \left(2\delta + \frac{2\delta^2(1+\delta)}{(1-\delta)} \right) = \frac{2\delta(1+\delta^2)}{1+\delta^2-2\delta} \leq \frac{2\delta}{1-2\delta}.$$

□

Proposition I.3 *The estimator $\widehat{G}$ studied in the previous proposition can be expressed as*

$$\widehat{G} = \sum_{\theta \in \Theta_\rho} \big[\widehat{\xi}_+(\theta) - \widehat{\xi}_-(\theta) \big] \theta\theta^\top$$

where

$$\begin{aligned}\big[\widehat{\xi}_+(\theta), \widehat{\xi}_-(\theta) \big]_{\theta \in \Theta_\rho} \in \arg \max_{(\xi_+, \xi_-) \in (\mathbb{R}_+^2)^{\Theta_\rho}} & - \frac{1}{2} \sum_{(\theta,\theta') \in \Theta_\rho^2} \big[\xi_+(\theta) - \xi_-(\theta) \big] \big[\xi_+(\theta') - \xi_-(\theta') \big] \langle \theta, \theta' \rangle^2 \\ & + \sum_{\theta \in \Theta_\rho} \xi_+(\theta) \widehat{N}(\theta)(1-\delta) - \xi_-(\theta) \widehat{N}(\theta)(1+\delta).\end{aligned}$$

Proof of Proposition I.3. Let us put

$$B_-(\theta) = \widehat{N}(\theta)(1-\delta) \quad \text{and} \quad B_+(\theta) = \widehat{N}(\theta)(1+\delta).$$

The estimated matrix $\widehat{G}$ is solution of the minmax optimisation problem

$$V = \inf_{H, H^\top = H} \sup_{(\xi_+, \xi_-) \in \left(\mathbb{R}_+^2\right)^\Theta} V(H, \xi_+, \xi_-),$$

where

$$V(H, \xi_+, \xi_-) = \frac{1}{2}\,\mathbf{T}\mathrm{r}(H^2) + \sum_{\theta \in \Theta} \xi_+(\theta)\big[B_-(\theta) - \theta^\top H \theta\big] + \xi_-(\theta)\big[\theta^\top H \theta - B_+(\theta)\big].$$

On the event Ω'' of probability at least $1-\epsilon$ the constraints are satisfied when $H = G$, and therefore, since they are linear constraints, Slater's conditions are satisfied [7, page 226]. This means that there is no duality gap, or in other words that

$$\inf_H \sup_{\xi_+, \xi_-} V(H, \xi_+, \xi_-) = \sup_{\xi_+, \xi_-} \inf_H V(H, \xi_+, \xi_-).$$

It is then elementary to compute explicitly the solution of

$$\inf_{H, H^\top = H} V(H, \xi_+, \xi_-),$$

that is

$$\widehat{H}(\xi_+, \xi_-) = \sum_{\theta \in \Theta_\rho} \big[\xi_+(\theta) - \xi_-(\theta)\big]\theta\theta^\top.$$

Remark that

$$\begin{aligned} V\big(\widehat{H}(\xi_+, \xi_-), \xi_+, \xi_-\big) \\ &= -\frac{1}{2} \sum_{(\theta, \theta') \in \Theta_\rho^2} \big[\xi_+(\theta) - \xi_-(\theta)\big]\big[\xi_+(\theta') - \xi_-(\theta')\big]\langle \theta, \theta' \rangle^2 \\ &\qquad + \sum_{\theta \in \Theta_\rho} \xi_+(\theta) B_-(\theta) - \xi_-(\theta) B_+(\theta). \end{aligned}$$

Since

$$\sup_{\xi_+, \xi_-} \inf_H V(H, \xi_+, \xi_-) = \sup_{\xi_+, \xi_-} V\big(\widehat{H}(\xi_+, \xi_-), \xi_+, \xi_-\big) = \inf_H \sup_{\xi_+, \xi_-} V(H, \xi_+, \xi_-),$$

we deduce that $\widehat{G} = \widehat{H}(\widehat{\xi_+}, \widehat{\xi_-})$, as stated in the proposition. □

Let us remark that, since $B_-(\theta) < B_+(\theta)$, the constraints cannot be reached at the same time for $\xi_+(\theta)$ and for $\xi_-(\theta)$, so that either $\xi_+(\theta) = 0$ or $\xi_-(\theta) = 0$, implying that ξ_+ and ξ_- are the positive and the negative parts of $\xi = \xi_+ - \xi_-$. We could thus also write $\widehat{G}$ as $\widehat{G} = \sum_{\theta\in\Theta} \widehat{\xi}(\theta)\theta\theta^\top$, where

$$\widehat{\xi} \in \arg\max_{\xi\in\mathbb{R}^\Theta} \left(-\frac{1}{2} \sum_{\theta,\theta'\in\Theta} \xi(\theta)\langle\theta,\theta'\rangle^2\xi(\theta') \right.$$
$$\left. + \sum_{\theta\in\Theta} \max\{\xi(\theta),0\}B_-(\theta) - \min\{\xi(\theta),0\}B_+(\theta) \right).$$

J. Computation of the robust Gram matrix estimator

J.1. Computation of the robust estimator in a fixed direction. In this section, we give some details on the computation of the estimator $\widehat{N}(\theta)$ used in Proposition 1.1 on page 7. In this discussion, θ is fixed.

Let us remark first that

$$r_\lambda(0) = -\lambda^{-1}\psi(\lambda) < 0, \text{ and}$$
$$\ell_\lambda(\theta) \stackrel{\text{def}}{=} \lim_{\alpha\to+\infty} r_\lambda(\alpha\theta) = \frac{\Big|\big\{i; \langle\theta, X_i\rangle \neq 0\big\}\Big| \log(2) - \Big|\big\{i; \langle\theta, X_i\rangle = 0\big\}\Big|\psi(\lambda)}{n\lambda},$$

so that these two quantities can easily be computed.

In the case when $\ell_\lambda(\lambda) \leq 0$, $\widehat{\alpha} = +\infty$ and $\widehat{N}(\theta) = 0$.

Otherwise, and this is the most usual case, $\ell_\lambda(\theta) > 0$, and we can compute $\widehat{\alpha}(\theta)$ and therefore $\widehat{N}(\theta)$ quickly using a modified Newton's method with global convergence properties.

Indeed, two algorithms may come to mind to compute $\widehat{\alpha}$.

The first is a divide and conquer algorithm, based on the fact that $\alpha \mapsto r_\lambda\big(\alpha\theta\big)$ is non-decreasing. Starting from $a_0, b_0 \in \mathbb{R}_+$ such that $r_\lambda\big(\sqrt{a_0}\theta\big) \leq 0$ and $r_\lambda\big(\sqrt{b_0}\theta\big) \geq 0$, we can put

$$a_k = \max\left\{u \in \left\{a_{k-1}, \frac{a_{k-1}+b_{k-1}}{2}, b_{k-1}\right\}, r\big(\sqrt{u}\theta\big) \leq 0\right\},$$
$$b_k = \min\left\{u \in \left\{a_{k-1}, \frac{a_{k-1}+b_{k-1}}{2}, b_{k-1}\right\}, r\big(\sqrt{u}\theta\big) > 0\right\}.$$

The second algorithm is the well known Newton's method, which is described in this case as

$$\alpha_k^2 = \alpha_{k-1}^2 - \frac{r_\lambda\big(\alpha_{k-1}\theta\big)}{\displaystyle\int \langle\theta, x\rangle^2 \psi'\big[\lambda\big(\alpha_{k-1}^2\langle\theta, x\rangle^2 - 1\big)\big]\, \mathrm{d}\overline{\mathbb{P}}(x)}.$$

Once we have reached a small enough neighborhood of the solution, Newton's method is faster, whereas the divide and conquer algorithm achieves a convergence speed of 2^{-k}, concerning the accuracy of the computation of $\widehat{N}(\theta)^{-1}$, from any starting point, not necessarily close to the solution.

Fortunately, it is quite easy to combine the two methods into a single algorithm that keeps the best of both worlds.

Let us introduce the function

$$f(u) = u - \frac{r_\lambda\big(\sqrt{u}\theta\big)}{\displaystyle\int \langle\theta, x\rangle^2 \psi'\big[\lambda\big(u\langle\theta, x\rangle^2 - 1\big)\big]\, \mathrm{d}\overline{\mathbb{P}}(x)}$$

Let us put to define the starting point of the algorithm

$$A_0 = \left\{0, \left(\int \langle\theta, x\rangle^2\, \mathrm{d}\overline{\mathbb{P}}(x)\right)^{-1}, \max\left\{\frac{1+\lambda^{-1}}{\langle\theta, X_i\rangle^2}, \langle\theta, X_i\rangle \neq 0, 1 \leq i \leq n\right\}\right\},$$

and

$$\begin{aligned} a_0 &= \max\big\{u \in A_0; r_\lambda\big(\sqrt{u}\theta\big) \leq 0\big\}, \\ b_0 &= \min\big\{u \in A_0; r_\lambda\big(\sqrt{u}\theta\big) > 0\big\}. \end{aligned}$$

Let us define then

$$A_k = \Big\{a_{k-1}, b_{k-1}, \frac{a_{k-1}+b_{k-1}}{2}, f(a_{k-1}), f(b_{k-1})\Big\},$$

$$\begin{aligned} a_k &= \max\big\{u \in A_k; r_\lambda\big(\sqrt{u}\theta\big) \leq 0\big\}, \\ b_k &= \min\big\{u \in A_k; r_\lambda\big(\sqrt{u}\theta\big) > 0\big\}, \end{aligned}$$

and $\alpha_k = \sqrt{a_k}$ if $r_\lambda\big(\sqrt{(a_k+b_k)/2}\,\theta\big) > 0$ and $\alpha_k = \sqrt{b_k}$ otherwise.

Proposition J.1

$$|\alpha_k^2 - \widehat{\alpha}^2| = \min\big\{|a_k - \widehat{\alpha}^2|, |b_k - \widehat{\alpha}^2|\big\} \leq 2^{-(k+1)}|b_0 - a_0|,$$

and

$$|\alpha_k^2 - \widehat{\alpha}^2| \leq \big|f\big(\alpha_{k-1}^2\big) - \widehat{\alpha}^2\big|.$$

The second inequality shows that the local convergence speed of the combined algorithm is at least as good as Newton's method.

J.2. Computation of a robust estimate of the Gram matrix. We describe here a simplified algorithm, that does not share the mathematical properties of the convex optimization scheme described in Appendix I, but turns out to be efficient in practice to improve on the empirical Gram matrix when dealing with the estimation of the Gram matrix, or the empirical risk minimization when dealing with least squares regression.

Since our approach of robust least squares regression in dimension d is based on the robust estimation of a Gram matrix in dimension $d+1$, we start with the robust estimation of the Gram matrix from a sample $X_1, \dots, X_n \in \mathbb{R}^d$ made of independent copies of some random variable X.

For any vector of weights $p = (p_i, i = 1, \dots, n) \in \mathbb{R}^n$, and any positive parameter λ, let $S(p, \lambda)$ be the solution of

$$\sum_{i=1}^n \psi\Big[\lambda\Big(S(p,\lambda)^{-1} p_i^2 - 1\Big)\Big] = 0,$$

that can be computed as explained in the previous section, using some suitable Newton algorithm. Given some confidence parameter ϵ, define $S(p) = S\big(p, \lambda(p)\big)$, where

$$\lambda(p) = m\sqrt{\frac{1}{v}\Big[\frac{2}{n}\log(\epsilon^{-1})\Big(1 - \frac{2}{n}\log(\epsilon^{-1})\Big)\Big]},$$
$$\text{with } m = \frac{1}{n}\sum_{i=1}^n p_i^2 \quad \text{and} \quad v = \frac{1}{n}\sum_{i=1}^n \big(p_i^2 - m\big)^2.$$

This value of the scale parameter is based on the optimal value for the estimation of a single expectation as described in [10] and maybe expected in practice to be more efficient than the conservative value allowing for the mathematical proof of generalization bounds.

The algorithm to compute a robust estimate $\widehat{G}$ of the Gram matrix $G = \mathbb{E}(XX^\top)$ works as follows.

Start with the empirical Gram matrix estimate

$$\widehat{G}_0 = \overline{G} = \frac{1}{n}\sum_{i=1}^n X_i\, X_i^\top.$$

Assuming that at iteration k we have computed the estimate $\widehat{G}(k) \in \mathbb{R}^{d\times d}$, decompose it into

$$\widehat{G}(k) = U(k)^\top D(k) U(k)$$

where $U(k)\, U(k)^\top = I$ is an orthogonal matrix and $D(k)$ is a diagonal matrix. Define the $d \times d$ matrix

$$M(k)_{i,j} = \frac{1}{4}\Bigg[S\Big((U(k)X_\ell)_i + (U(k)X_\ell)_j, \ell = 1, \dots, n\Big) - S\Big((U(k)X_\ell)_i - (U(k)X_\ell)_j, \ell = 1, \dots, n\Big)\Bigg],$$

and update the estimate by the formula

$$\widehat{G}(k+1) = U(k)^\top M(k)\, U(k),$$

until some stopping rule is reached (we can for instance use a fixed number of iterations, as in the simulation below, or stop when the Frobenius norm $\|\widehat{G}(k+1) - \widehat{G}(k)\|_{\mathrm{F}}$ falls under some threshold). The idea is to update the estimator of the Gram matrix using the polarization formula (1) on page 3 in a basis of eigenvectors of the current estimate. This uses more directions than using the polarization formula only in the canonical basis of $\mathbb{R}^d$, while trying to get accurate eigenvectors, and is faster than using a net of directions as in the mathematically more justified algorithm described in Appendix I.

To solve the least squares problem

$$\inf_{\theta \in \mathbb{R}^d} \mathbb{E}\big[\big(Y - \langle \theta, X\rangle\big)^2\big]$$

from a sample $(X_1, Y_1), \dots, (X_n, Y_n)$ of independent copies of $(X, Y) \in \mathbb{R}^{d+1}$, first compute as above a robust estimate

$$\widehat{G} = \begin{pmatrix} \widehat{G}_{1,1} & \widehat{G}_{1,2} \\ \widehat{G}_{2,1} & \widehat{G}_{2,2} \end{pmatrix}$$

of the Gram matrix

$$\mathbb{E}\left[\begin{pmatrix} X \\ -Y \end{pmatrix}\big(X^\top, -Y\big)\right],$$

then define the robust estimate $\widehat{\theta}$ of θ_* as

$$\widehat{\theta} = -\widehat{G}_{1,1}^{-1}\widehat{G}_{1,2},$$

where $\widehat{G}_{1,1}^{-1}$ is the pseudo inverse of the symmetric matrix $\widehat{G}_{1,1}$, obtained by inverting only its non-zero eigenvalues.

K. Some simulation

We present a small simulation to illustrate the benefit of using our robust least squares estimator in the case of a long tail noise. The idea of this

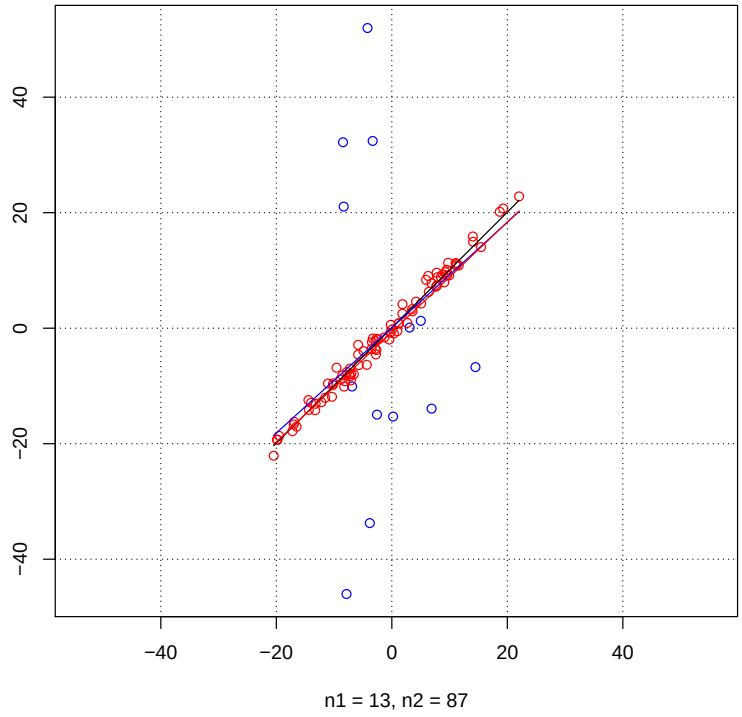

Figure 1: A sample

simulation is to show that even in a very simple situation using a robust estimator may bring a significant improvement.

Consider some noise $\eta \sim 0.9 \times \mathcal{N}(0,1) + 0.1 \times \mathcal{N}(0, 30^2)$ that is the mixture of two Gaussian random variables with different variances. Consider a Gaussian random variable $\widetilde{X} \sim \mathcal{N}(0, 10^2)$ independent of the noise η, and define

$$Y = \widetilde{X} + \eta + 1.$$

Putting $X = (\widetilde{X}, 1)^\top$, we can write this problem as

$$Y = \langle \theta_*, X \rangle + \eta,$$

where $\theta_* = (1,1)^\top$.

Figure 1 shows a typical sample, where $n = 100$, and where the two components of the mixture have been plotted with different colors. Figures 2 on the following page plots the excess risk of the empirical risk minimizer

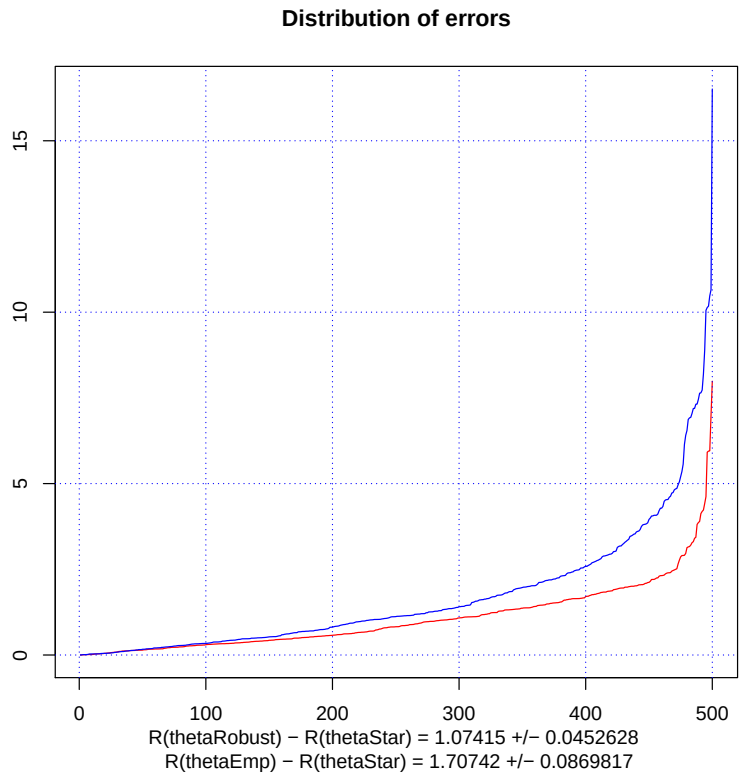

Figure 2: The empirical quantile function of the excess risk

and of the robust estimate for 500 trials of the experiment. One can see that we get a substantial improvement of the mean excess risk using the robust estimator, since the estimated expected excess risks are of 1.7 against less than 1.1 for the robust estimator.

In conclusion, it is not necessary to envision very large sample sizes or very exotic noise structures to feel the improvement brought by the more stable robust estimator.

References

[1] P. Alquier. Iterative feature selection in least square regression estimation. *Ann. Inst. Henri Poincaré, Probab. Stat.*, 2008.

[2] P. Alquier. PAC-bayesian bounds for randomized empirical risk minimizers. *Mathematical Methods of Statistics*, 17(4):279–304, 2008.

[3] J.-Y. Audibert. Aggregated estimators and empirical complexity for least square regression. *Ann. Inst. Henri Poincaré, Probab. Stat.*, 40(6):685–736, 2004.

[4] J.-Y. Audibert and O. Catoni. Linear regression through PAC-bayesian truncation. *arXiv* http://arxiv.org/abs/1010.0072v2, pages 1–40, 2010, revised in 2011.

[5] J.-Y. Audibert and O. Catoni. Robust linear least squares regression. *Ann. Stat.*, 39(5):2766–2794, 2011.

[6] J.-Y. Audibert and O. Catoni. Supplement to "robust linear least squares regression.". *Ann. Stat.*, pages 1–19, 2011.

[7] S. Boyd and L. Vandengerghe. *Convex Optimization*. Cambridge University Press, 2004.

[8] O. Catoni. *Statistical Learning Theory and Stochastic Optimization, Lectures on Probability Theory and Statistics, École d'Été de Probabilités de Saint-Flour XXXI – 2001*, volume 1851 of *Lecture Notes in Mathematics*. Springer, 2004. Pages 1–269.

[9] O. Catoni. *PAC-Bayesian Supervised Classification: The Thermodynamics of Statistical Learning*, volume 56 of *IMS Lecture Notes Monograph Series*. Institute of Mathematical Statistics, 2007. Pages i-xii, 1-163.

[10] O. Catoni. Challenging the empirical mean and empirical variance: a deviation study. *Ann. Inst. Henri Poincaré*, 48(4):1148–1185, 2012.

[11] O. Catoni. Pac-bayes bounds for supervised classification. In A. Gammerman V. Vovk, H. Papadopoulos, editor, *Measures of Complexity, Festschrift for Alexey Chervonenkis*, pages 287–302. Springer, 2015.

[12] I. Giulini. *Generalization bounds for random samples in Hilbert spaces*. PhD thesis, Ecole Normale Supérieure, Paris, France, 2015.

[13] I. Giulini. Robust dimension-free gram operator estimates. preprint, 2015. http://arxiv.org/abs/1511.06259.

[14] V. Koltchinskii and S. Mendelson. Bounding the smallest singular value of a random matrix without concentration. *Int Math Res Notices*, 2015.

[15] G. Lecué and S. Mendelson. Performance of empirical risk minimization in linear aggregation. *Bernoulli to appear*, 2015.

[16] D. A. McAllester. PAC-Bayesian model averaging. In *Proceedings of the 12th annual conference on Computational Learning Theory*. Morgan Kaufmann, 1999.

[17] D. A. McAllester. PAC-Bayesian stochastic model selection. *Mach. Learn.*, 51(1):5–21, April 2003.

[18] D. A. McAllester. Simplified PAC-Bayesian margin bounds. In *In COLT*, pages 203–215, 2003.

[19] M. Seeger. PAC-Bayesian generalization error bounds for gaussian process classification. Informatics report series EDI-INF-RR-0094, Division of Informatics, University of Edinburgh, 2002.